



Intel[®] C/C++ Compiler User's Guide

With Support for the Streaming SIMD Extensions

Copyright © 1996–1999 Intel Corporation
All Rights Reserved
Issued in U.S.A.
Order Number 718195-001

Intel[®] C/C++ Compiler User's Guide: With Support for the Streaming SIMD Extensions

Order Number: 718195-001

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

This Intel® C/C++ Compiler User's Guide as well as the software described in it is furnished under license and may only be used or copied in accordance with the terms of the license. The information in this manual is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by Intel Corporation. Intel Corporation assumes no responsibility or liability for any errors or inaccuracies that may appear in this document or any software that may be provided in association with this document.

Except as permitted by such license, no part of this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the express written consent of Intel Corporation.

Intel may make changes to specifications and product descriptions at any time, without notice.

* Third-party brands and names are the property of their respective owners.

Copyright © Intel Corporation 1996–1999. Third-party brands and names are the property of their respective owners.

Contents

About This Manual

Related Publications	xiv
Notational Conventions	xv

Chapter 1 Overview

Tools You Need	1-1
Application Development	1-2

Chapter 2 Compiler Operation

Compiler Command-line Syntax	2-1
Using the Intel® C/C++ Compiler within the Microsoft Visual C++ Integrated Development Environment (IDE)	2-2
Filename Extensions	2-3
Compiler Option Quick Guide	2-4
Default Behavior of the Compiler	2-15

Chapter 3 Changing the Compilation Environment

Environment Variables	3-1
Configuration Files	3-2
Response Files	3-2
Include Files	3-3
Specifying an Include Directory (-I)	3-3
Removing Include Directories (-X)	3-3

How To Specify Alternate Tools and Paths	3-4
Specifying an Alternate Component (-Qlocation,tool,path)	3-4
Passing Options to Other Programs (-Qoption, tool, optlist)	3-4
Passing Options to the Linker (-link)	3-5
Chapter 4 Optimizations	
Optimization Choices	4-1
Restricting Optimization with -Od	4-2
The Effect of -Od on -Oy and Debugging	4-2
Targeting a Processor (-Gn)	4-3
Automatic Processor Dispatch Support (-Qx[extensions] and -Qax[extensions])	4-3
Exclusive Specialized Code with -Qx[extensions]	4-4
Specialized and Generic Code with -Qax[extensions] ..	4-4
Prefetching with -Qpf[options]	4-5
Loop Unrolling with -Qunrolln	4-6
Inline Expansion of Library Functions (-Oi, -Oi-)	4-6
Floating-point Arithmetic Precision (-Op, -Op-, -Qprec, -Qprec_div, -Qpc, -Qlong_double)	4-7
When to Use -Op	4-7
When to Use -Qprec	4-8
When to Use -Qprec_div	4-8
When to Use -Qpc	4-8
Alternatives to the -Qpc Option	4-9
When to Use -Qlong_double	4-10
Rounding Control Option (-Qrcd)	4-10
Chapter 5 Interprocedural and Profile Guided Optimizations	
Interprocedural Optimization (IPO) with -Qip and -Qipo	5-2
Multifile IPO (-Qipo)	5-2

Creating a Multifile IPO Executable.....	5-3
Creating a Multifile IPO Executable	
Using a Project Makefile	5-4
Controlling Inline Expansion of User Functions	
(-Obn, -Qip_no_inlining).....	5-4
Profile-Guided Optimization (PGO): Three Phases	5-5
Basic PGO Options and Environment Variables.....	5-7
Using Profile-Guided Optimization: An Example	5-8
Guidelines for Using PGO	5-8
Profile Guided Optimizations Using	
a Function Order List.....	5-9
Function Order List Usage Guidelines.....	5-10
Function Order List Example	5-10
Utilities for Profile-Guided Optimization.....	5-11
The profmerge Utility	5-11
The proforder Utility	5-11
Function Call to Dump Profile Data Explicitly	5-12

Chapter 6 Specifying Compilation Output

Parsing for Syntax Only (-Zs).....	6-2
Producing an Assembly Code Listing (-S)	6-2
Suppressing Linking (-c)	6-4
Using the Microsoft Assembler to Produce	
Object Code (-Quse_asm).....	6-4
Linking.....	6-5
Naming the Output File (-Fe, -Fo, -Fa).....	6-5
Preparing for Debugging (-Zi, -Oy, -Oy-).....	6-6
Optimizations and Support for Symbolic Debugging.....	6-7

Chapter 7 Preprocessing

Preserving Comments in	
Preprocessed Source Output (-C)	7-2
Preprocessing Only (-E, -EP, and -P).....	7-2

Defining Macros (-QA, -QA-, -u, -D, and -U).....	7-3
Predefined Macros.....	7-4
Printing Include-file Dependencies (-QH)	7-5
Printing Makefile Dependencies (-QM).....	7-6
 Chapter 8 C/C++ Language Features	
Conformance to C Standards	8-1
C Language Dialects.....	8-2
Strict ANSI Dialect (-Za)	8-2
Extended Dialect (-Ze)	8-2
Predefined Macros for Standard Conformance	8-4
Conformance to C++ Standards	8-5
 Chapter 9 Microsoft Compatibility	
Compiler Pragmas	9-1
Microsoft Compatibility Option (-Qms).....	9-2
Microsoft Version Compatibility Options (-Qvcn)	9-2
Unsupported Compiler Options	9-3
Differences in PCH Support.....	9-4
Compilation and Execution Differences.....	9-4
Inline Assembly Target Labels	9-4
Preprocessor Macro Expansion.....	9-5
Evaluation of Left Shift Operations	9-6
Use of Friend Injection: Not Recommended.....	9-7
Declaration in Scope of Function Defined in a Namespace ..	9-8
Enum Bit-Field Signedness	9-8
Debugging Functions with Aligned Stack Frames	
Using MSVC++ 4.2.....	9-8
 Chapter 10 Diagnostic Information	
Disabling the Sign-on Message (-nologo).....	10-1
Printing the List of icl Options (-?, -help).....	10-1
Diagnostic Messages.....	10-2

Suppressing Warning Messages with lint Comments	10-4
Suppressing Warning Messages or Enabling Remarks (-w, -Wn)	10-4
Controlling the Severity of Diagnostics (-Qwd, -Qwr, -Qww, -Qwe)	10-5
Limiting the Number of Errors Reported (-Qwnum)	10-6
Additional Information about the Compilation	10-6
Chapter 11 Libraries	
Managing Libraries	11-1
Default Libraries	11-2
Library Files	11-2
Math Libraries	11-3
Enabling the Floating-Point Division Check (-Qlfddiv)	11-3
Avoiding Incorrect Decoding of Certain Instructions (-QIOf)	11-4
Chapter 12 Controlling Compiler-Generated Code	
Specifying Structure Tag Alignments (-Zp)	12-1
Allocation of Zero-initialized Variables (-Qnobss_init)	12-2
Chapter 13 Support for MMX™ Technology and the Streaming SIMD Extensions	
MMX Technology Intrinsics	13-1
The EMMS Instruction: Why You Need it and When to Use it	13-2
Guidelines for When to Use EMMS	13-3
MMX Technology Intrinsic Groups	13-4
General Support Intrinsics	13-5
Packed Arithmetic Intrinsics	13-7
Shift Intrinsics	13-9
Logical Intrinsics	13-12
Compare Intrinsics	13-13

Processor-Dispatching Support	13-14
Streaming SIMD Extensions Intrinsics	13-17
The Intrinsics API	13-17
The __m128 Data Type	13-18
Streaming SIMD Extensions Intrinsic Conventions ..	13-18
Floating Point Intrinsics	13-19
Arithmetic Operations	13-20
Logical Operations	13-23
Comparisons	13-24
Conversion Operations	13-31
Miscellaneous	13-34
Macro Function for Shuffle	13-36
Macro Functions to Read and Write the Control Registers	13-37
Macro Function for Matrix Transposition	13-39
Memory and Initialization	13-40
Load Operations	13-40
Set Operations	13-41
Store Operations	13-42
Integer Intrinsics	13-43
Cacheability Support	13-47
Data Alignment	13-47
Alignment Support	13-48
Dynamic Stack Frame Alignment	13-49
Assembly Language Support	13-50
Inline Assembly	13-50
Generation of Assembly files	13-51

Chapter 14 Compiler Vectorization Support and Guidelines

Vectorizer Quick Reference	14-1
Command-Line Switch Support	14-2
Language Support and Pragmas	14-5

Alignment with declspec	14-6
Qualifying with the restrict Keyword	14-6
Pragma Scope	14-7
Loop Structure Coding Background	14-10
Key Programming Guidelines	14-10
Loop Constructs	14-11
Loop Body Control Flow	14-11
Loop Exit Conditions	14-12
Stripmining and Cleanup	14-14
Statements in the Loop Body	14-14
Floating-point array operations	14-14
Integer Array Operations	14-15
Other Integer Operations	14-15
Other Datatypes	14-15
No Function Calls	14-15
Vectorizable Data References	14-15
Data alignment	14-16
Common Errors in Making Code	
Vectorization-Compatible	14-16
Some Examples	14-17
Argument Aliasing: A Vector Copy	14-17
Data Alignment: Two Examples	14-18
Data Alignment Examples	14-18
Data Dependence	14-20
Loop Interchange and Subscripts: Matrix Multiply	14-22
For Additional Information	14-23

Appendix A Compiler Limits

Appendix B Experimental Performance Tuning

Keywords for Optimization (-Qoption,c,optlist)	B-1
Interprocedural Optimization.....	B-2
Analyzing the Effects of Multifile	
IPO (-Qipo_c, -Qipo_S)	B-3
Using In-line Heuristics (-Qinl_heur n)	B-3
Criteria for In-Line Function Expansion	B-4

Index

Examples

3-1 Sample icl.cfg File	3-2
13-1 Correct EMMS Usage In Initialization Code	13-4
13-2 Shuffle Function Macro	13-32
13-3 View of Original and Result Words With Shuffle Function Macro 1	3-32
13-4 Exception State Macros with _MM_EXCEPT_DIV_ZERO.....	13-33
13-5 Exception Mask with _MM_MASK_OVERFLOW and _MM_MASK_UNDERFLOW	13-33
13-6 Rounding Mode with _MM_ROUND_TOWARD_ZERO	13-34
13-7 Flush-to-Zero Mode with _MM_FLUSH_ZERO_OFF	13-34
14-1 Using -Qvec_no_arg_alias to Vectorize	14-4
14-2 Using -Qvec_no_alias to Vectorize	14-5
14-3 Using declspec(align(n)) for Alignment	14-6
14-4 Using the restrict Keyword as a Qualifier	14-7
14-5 Loop Using #pragma ivdep	14-8
14-6 Loop Using Stride-2 #pragma vector	14-8
14-7 Loop Using #pragma vector aligned	14-9
14-8 Low Trip Count Loop Using #pragma novector	14-9

14-9 Loop Construct Usage	14-11
14-10 Loop Body Control Flow	14-12
14-11 Loop Usage Comparisons	14-13
14-12 Strip Mining and Cleanup Loops	14-14
14-13 Vectorizable Loop Invariant Reference	14-16
14-14 Unvectorizable Copy Due to Unproven Distinction .	14-17
14-15 Using restrict to Prove Vectorizable Distinction	14-17
14-16 Unaligned Loop Due to Global Variables	14-18
14-17 Loop Alignment with declspec	14-18
14-18 Loop Unaligned Due to Unknown Variable Value at Compile Time	14-19
14-19 Alignment Due to Assertion of Variable as Multiple of 4	14-19
14-20 Data Dependent Loop	14-20
14-21 Data Dependence Vectorization Patterns	14-21
14-22 Data Independent but Still Unvectorizable Loop	14-22
14-23 Typical Matrix Multiplication	14-22
14-24 Matrix Multiplication with Stride-1	14-23

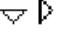



















Figures

1-1 Application Development Cycle	1-2
4-1 Example of the _controlp() Function:	4-10
5-1 Phases of Basic Profile Guided Optimization	5-6
9-1 Simplified View of Arbitrary Offset in the Stack	9-9
13-1 Why You Need EMMS to Reset After an MMX™ Instruction	13-2
13-2 Matrix Transposition Using the _MM_TRANSPOSE4_PS Macro	13-35
14-1 A Redistribution for Data Dependence	14-20

Tables

2-1	Default Filename Extensions	2-3
2-2	Summary of Command-line Options	2-4
4-1	Optimization Summary	4-2
4-2	Processor Dispatch Extension Options	4-4
4-3	-Qpc Equivalent Calls	4-9
5-1	IPO Optimization Summary	5-2
5-2	Summary of -Obn and -Qip_no_inlining Options	5-4
5-3	Basic Profile-Guided Optimization Options	5-7
5-4	Profile-Guided Optimization Environment Variables	5-7
5-5	Advanced Profile-Guided Optimization Options	5-9
6-1	Compiler Input and Output Summary	6-1
6-2	Effects of Using -Zi with Optimization Options	6-8
7-1	Options to Control Preprocessing	7-1
7-2	Predefined Macros	7-4
8-1	Predefined Macros for Standard Conformance	8-5
9-1	Microsoft Version Compatibility Options	9-2
9-2	List of Unsupported Microsoft Visual C++ Compiler Options	9-3
13-1	General Support Intrinsics	13-5
13-2	Packed Arithmetic Intrinsics	13-7
13-3	Shift Intrinsics	13-9
13-4	Logical Intrinsics	13-12
13-5	Compare Intrinsics	13-13
13-6	Conversion Operators	13-31
13-7	Integer Intrinsics	13-43
13-8	Stack Frame Alignment Options	13-50
14-1	Vectorization Command-Line Switches	14-1
14-2	Language Support	14-5
A-1	Compiler Limits	A-1

How to Use This Online Manual

	Click to hide or show subtopics when the bookmarks are shown.		Click to go to the previous page.
	Double-click to jump to a topic when the bookmarks are shown.		Click to go to the next page.
	Click to display bookmarks.		Click to go to the last page.
	Click to display thumbnails.		Click to return to the previous view. Use this button when you need to go back after using the jump button (see below).
	Click to close bookmark or thumbnail view.		Click to go forward from the previous view.
	Click and use on the page to drag the page in vertical direction.		Click to set 100% of the page view.
	Click and drag on the page to magnify the view.		Click to display the entire page within the window.
	Click and drag on the page to reduce the view.		Click to fill the width of the window.
	Click and drag the selection cursor to the page.		Click to open a dialog to search for a word or multiple words.
	Click to go to the first page of the manual.		Click jump button on manual pages to jump to the related subjects. Use the return back icon above to go back.

Printing an Online File: Select **Print** from the **File** menu to print an online file. The dialog that opens allows you to print full text, range of pages, or selection.

Viewing Multiple Online Manuals: Select **Open** from the **File** menu, and open the .PDF file you need. Select **Cascade** from the **Window** menu to view multiple files.

Resizing the Bookmark Area: Drag the double-headed arrow that appears on the area's border as you pass over it.

Jumping to Topics: In this manual, many of the function names and section titles are printed in green color, to indicate that you can jump to them.

On the next page, you will find a jump text demo.

Jump Text Demo

The following text segment is taken from the manual. In this example, the section name “Restricting Optimization” is displayed in green underlined font. If you click this section name, the first page of the section will display. To return to this page

click the  icon in the tool bar.

This chapter describes ways to improve the performance of your application and the effects of optimization options on programs. The [“Restricting Optimization with -Od”](#) section tells you how to suspend optimizations for certain applications or for debugging.

About This Manual

This manual describes how to use the Intel[®] C/C++ Compiler for Win32* systems. The Intel C/C++ Compiler can be hosted on either a Windows NT* or Windows* 95 operating system.

- Chapters 1 through 3 help you get started.
- The [“Compiler Option Quick Guide” in Chapter 2](#) alphabetically lists and briefly describes compiler options. It also provides cross-reference links to more detailed descriptions in this manual.
- Chapter 4 describes well-known performance optimizations.
- Chapter 5 describes Interprocedural and Profile-Guided Optimizations.
- Chapters 6 through 9 help you use preprocessing and maintain language conformance.
- Chapter 10 describes diagnostic information.
- Chapter 11 describes libraries.
- Chapter 12 describes how to control compiler-generated code.
- Chapter 13 describes support for MMX[™] Technology and the Streaming SIMD Extensions.
- Chapter 14 describes vectorization options.
- Appendix A provides a table of the compiler limits.
- Appendix B provides information on experimental performance tuning.
- This manual also provides a glossary and an index.

This manual does not teach you the C programming languages. As prerequisites, you should be familiar with software development of C/C++ and assembly-language programs. You should also be familiar with the host computer’s operating system and architecture of Intel[®] processors.

The Intel C/C++ Compiler operates similarly to the Microsoft Visual C++* compiler. In many cases, you can apply the functional descriptions provided with the Microsoft Visual C++ compiler to the Intel compiler, except where this manual states that you cannot. See [“Application Development” in Chapter 1](#) for the relationship of the Intel C/C++ Compiler to system-specific programming support tools.

Related Publications

The following documents provide additional information relevant to the Intel C/C++ Compiler:

- *The Annotated C++ Reference Manual*, 3rd edition, Ellis, Margaret; Stroustrup, Bjarne, Addison Wesley, 1991. Provides information on the C++ programming language.
- *The C Programming Language*, 2nd edition, Kernighan, Brian W.; Ritchie, Dennis W., Prentice Hall, 1988. Provides information on the K & R definition of the C language.
- *C: A Reference Manual*, 3rd edition, Harbison, Samuel P.; Steele, Guy L., Prentice Hall, 1991. Provides information on the ANSI standard and extensions of the C language.
- For environment specifics, see the manuals or on-line help supplied with Microsoft Visual C++, 32-bit edition for Windows.
- For Win32 specific information, see the documentation included with the *Microsoft Win32 Software Development Kit, Version 3.1*.

Information about the target architecture is available from Intel and from most technical bookstores. Some helpful titles are:

- *Intel Architecture Software Developer's Manual, Volume 1: Basic Architecture*, Intel Corporation, order number 243190
- *Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference Manual*, Intel Corporation, order number 243190
- *Intel Processor Identification with the CPUID Instruction*, Intel Corporation, order number 241618
- *Intel Architecture MMX Technology Programmer's Reference Manual*, Intel Corporation, order number 241618.
- *Pentium® Pro Processor Developer's Manual (3 Volume Set)*. Intel Corporation, order number 242693.

-
- *Pentium II Processor Developer's Manual*. Intel Corporation, order number 243502-001.
 - *Pentium Processor Specification Update*. Intel Corporation, order number 242480.
 - *Pentium Processor Family Developer's Manual*, Intel Corporation, order numbers 241428-005.
 - Most Intel documents are also available from the Intel Corporation web site at www.intel.com.

Notational Conventions

This manual uses the following conventions:

<i>This type style</i>	Indicates an element of syntax, a reserved word, a keyword, a filename, computer output, or part of a program example. The text appears in lowercase unless uppercase is significant. <i>1</i> indicates lowercase letter L in examples. <i>1</i> is the number 1 in examples. <i>O</i> is the uppercase O in examples. <i>0</i> is the number 0 in examples.
This type style	Indicates the exact characters you type as input.
<i>This type style</i>	Indicates a placeholder for an identifier, an expression, a string, a symbol, or a value. Substitute one of these items for the placeholder.
[<i>items</i>]	Indicates that the items enclosed in brackets are options.
{ <i>item</i> <i>item</i> }	Indicates to elect only one of the items listed between braces. A vertical bar () separates the items.
... (ellipses)	Indicate that you can repeat the preceding item.

Overview

1

This chapter introduces you to the Intel C/C++ Compiler for Win32 Systems. The compiler is designed to operate on any 32-bit Intel Architecture machine. However, it produces optimized code for the Intel486™, Pentium®, Pentium Pro, Pentium II, Celeron™, Pentium II Xeon™, and Pentium III processors.

Tools You Need

To compile programs with this compiler you need the following tools:

- Windows NT, Windows 95, or Windows 98 operating systems
- Microsoft Visual C++ Versions 4.2 or higher available on the host operating system



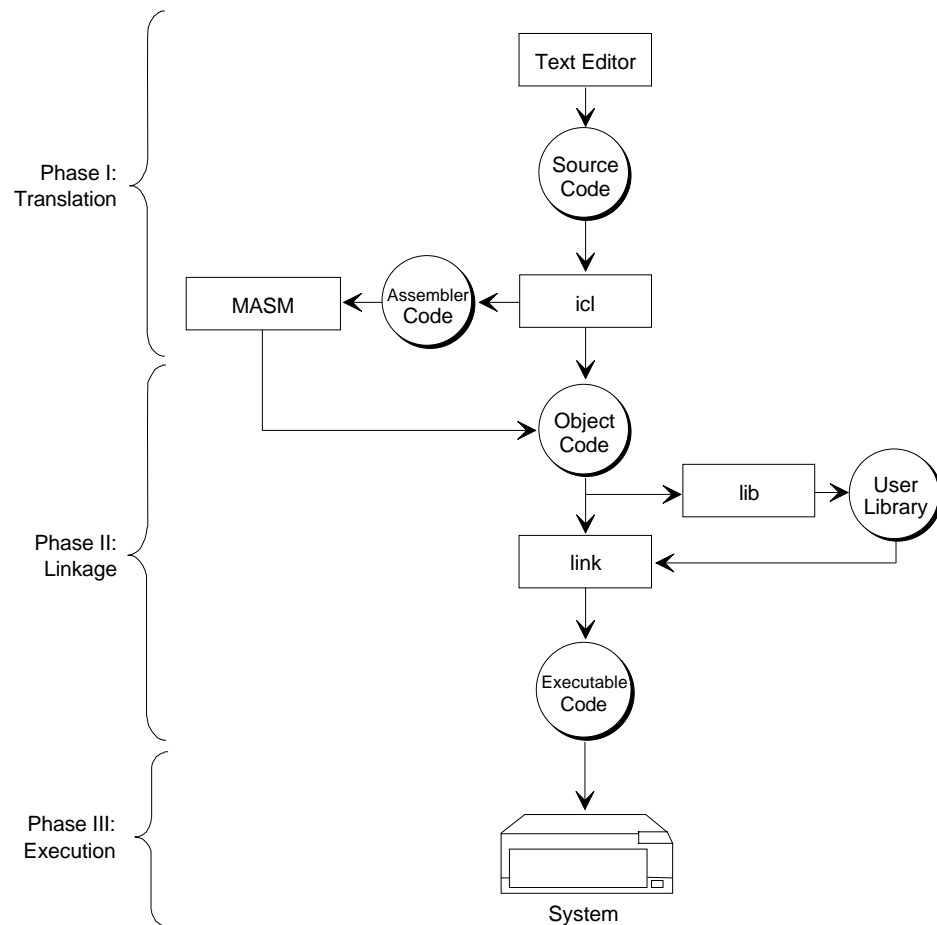
NOTE. *The Intel C/C++ Compiler is a plug-in product that uses the Microsoft Visual C++ header files, libraries, linker, and many other components.*

If you plan to assemble files generated by the assembler, you must provide an assembler. Version 6.12 of the Microsoft Macro Assembler (MASM) is recommended, but you can use Version 6.11 if you are not going to use the Streaming SIMD Extensions. Whatever assembler you use, you must name it **ML.EXE** and place it in your default path.

Application Development

Figure 1-1 shows you how the Intel compiler (**icl**) fits into the Microsoft application development environment. The compiler processes C or C++ language source and generates either assembly or object modules. You decide the input and output by setting options when you run the compiler.

Figure 1-1 Application Development Cycle



OSD2098

Compiler Operation

2

This chapter describes how to run the compiler, how to select compilation options, and the default behavior of the compiler.

This chapter also contains the [“Compiler Option Quick Guide”](#), which lists all the Intel C/C++ Compiler options. The options that you invoke change the compiler’s default behavior in the following areas:

- controlling compilation flow
- controlling compiler output
- preprocessing files
- controlling optimization
- specifying language conformance
- preparing for debugging
- controlling the linker
- managing special cases
- preparing for profiling

The options available for the Intel C/C++ Compiler are described in detail in the remaining chapters of this manual.

Compiler Command-line Syntax

To invoke the compiler from the DOS command line, enter the following command:

```
prompt> icl [options] [path]filenames
```

<i>options</i>	Indicates one or more command-line options. The compiler recognizes one or more letters preceded by a hyphen (-) or a forward slash (/) as an option.
<i>filenames</i>	Indicates one or more source files to be processed by the compilation system. You can specify more than one <i>[path]filename</i> . You must separate multiple filenames with a space.

Using the Intel® C/C++ Compiler within the Microsoft Visual C++ Integrated Development Environment (IDE)

If you have Microsoft Visual C++ versions 4.2 or higher on your system, the installation procedure automatically plugs the Intel C/C++ Compiler into the Microsoft IDE. This provides compatibility between the Intel and Microsoft C/C++ compiler switch options. Additionally, you are enabled with Intel's optimizations, which are designed specifically to exploit the power of Intel architectures. To invoke the Intel C/C++ Compiler from the Microsoft IDE, follow these steps:

1. Click on Select Compiler from the Tools Menu.
2. The Microsoft Developer Studio displays the Select Compiler window, which allows you to choose a compiler.
3. To choose the Intel Compiler, click the check-box and specify the version in the down-arrow menu.
4. Click the OK button to return to the Microsoft IDE.



NOTE. *To switch back to the Microsoft compiler, clear the box for the Intel C/C++ Compiler.*

You can now use Microsoft IDE to build your application or portions of your application using the selected compiler. For more information on using the Intel C/C++ Compiler, click the Help button from the Intel Compiler Selection Tool in the Visual C++ development environment.

Filename Extensions

By default the compiler recognizes `.cc`, `.cpp`, and `.cxx` files as C++ files. In examples, this guide uses the `.cpp` extension for C++ files. The compiler recognizes files with the `.c` extension as C files. Also, the Intel C/C++ Compiler recognizes the default filename extensions listed in [Table 2-1](#).

Table 2-1 **Default Filename Extensions**

Filename	Interpretation	Action
<code>filename.lib</code>	object library	Passed to LINK.exe.
<code>filename.i</code>	C or C++ source preprocessed and expanded by the C/C++ preprocessor	Passed to compiler.
<code>filename.obj</code>	compiled object module	Passed to LINK.exe.
<code>filename.asm</code>	assembly file	Assembled by MASM.

If you use specify `-TP [files]`, the compiler treats all files on the command line as C/C++ files. Both `x.i` and `y.c` are treated as C++ files on the following command line:

```
prompt> icl -c -TP x.i y.c
```

However, if you specify `-Tp file`, the compiler treats the next file on the command line as a C++ file regardless of the extension. For example, to treat `x.i` as a C++ file, but `y.c` as a C file, you could type the following:

```
prompt> icl -c y.c -Tp x.i
```



NOTE. *Usually, you cannot associate a filename with an option by command-line placement alone. Specified options apply to all files. For example, in the following command line, the `-Za` option applies to both `x.cpp` and `y.cpp`:*

```
prompt> icl -c x.cpp -Za y.cpp
```

The `-Tc` and `-Tp` are exceptions because they affect only the file that follows them on the command line.

Compiler Option Quick Guide

[Table 2-2](#) provides you with a reference of all compilation control options and some options for control of the linker.

Table 2-2 Summary of Command-line Options

Option	Description	Default	Reference
<code>-?, -help</code>	Prints compiler options summary.	OFF	page 10-1
<code>-c</code>	Stops the compilation process after an object file has been generated. The compiler generates an object file for each C or C++ source file or preprocessed source file.	OFF	page 6-4
<code>-C</code>	Places comments in preprocessed source output.	OFF	page 7-2
<code>-Dname [=value]</code>	Defines a macro name and associates it with the specified value.	OFF	page 7-3
<code>-E</code>	Stops the compilation process after the C or C++ source files have been preprocessed, and writes the results to <code>stdout</code> .	OFF	page 7-2
<code>-EHa</code>	Enables asynchronous C++ exception handling model.	OFF	*
<code>-EHc</code>	Specifies that <code>extern C</code> functions do not throw exceptions.	ON	*
<code>-EHs</code>	Enables the synchronous C++ exception handling model.	OFF	*
<code>-EP</code>	Stops the compilation process after the C or C++ source files have been preprocessed and writes the results to <code>stdout</code> . The <code>#line</code> directives are stripped (that is, not included in the output).	OFF	page 7-2
<code>-F n</code>	Sets the amount of stack to reserve for the program by passing <code>-stack:n</code> to the linker.	OFF	*

* The options marked with the asterisk in the "Reference" column produce the same results for the Intel C/C++ Compiler as for the Microsoft Visual C++ Compiler. If you need more information on these options, see the documentation that accompanies Microsoft Visual C++.

Table 2-2 Summary of Command-line Options (Continued)

Option	Description	Default	Reference
<code>-FA[c s]</code>	Produces an assembly output file. The <code>s</code> argument has no effect on the Intel C Compiler.	OFF	*
<code>-Fa[filename]</code>	Produces an assembly output file with the specified filename, or the default name if the filename is not specified.	OFF	page 6-5
<code>-Fe[filename]</code>	Produces an executable output file with the specified filename, or the default name if the filename is not specified.	ON	page 6-5
<code>-Fm[filename]</code>	Instructs the linker to produce a map file.	OFF	*
<code>-Fo[filename]</code>	Produces an object output file with the specified filename, or the default name if the filename is not specified.	ON	page 6-5
<code>-G4</code>	Targets the optimizations to the Intel i486 processor.	OFF	page 4-3
<code>-G5</code>	Targets the optimizations to the Intel Pentium processor.	OFF	page 4-3
<code>-G6</code>	Targets the optimizations to the Intel Pentium Pro and Pentium II processors.	ON	page 4-3
<code>-GA</code>	Optimizes for Microsoft Windows applications.	OFF	*
<code>-Ge</code>	Enables stack-checking.	ON	*
<code>-Gf</code>	Enables string-pooling optimization.	ON	*
<code>-GF</code>	Enables read-only string-pooling.	ON	*
<code>-Gh</code>	Adds a call to user-supplied <code>__penter</code> routine in each function prolog.	OFF	*
<code>-GR</code>	Enables C++ Runtime Type Information (RTTI).	OFF	*

* The options marked with the asterisk in the "Reference" column produce the same results for the Intel C/C++ Compiler as for the Microsoft Visual C++ Compiler. If you need more information on these options, see the documentation that accompanies Microsoft Visual C++.

Table 2-2 Summary of Command-line Options (Continued)

Option	Description	Default	Reference
<code>-Gr</code>	Selects Microsoft <code>_fastcall</code> as the default calling convention.	OFF	*
<code>-GR-</code>	Disables C++ RTTI.	ON	*
<code>-Gs n</code>	Enables stack-checking for routines with <code>n</code> or more bytes of local variables and compiler temporaries.	OFF	*
<code>-GT</code>	Enables fiber-safe thread local storage (TLS).	OFF	*
<code>-GX</code>	Enables C++ exception handling.	OFF	*
<code>-GX-</code>	Disables C++ exception handling.	ON	*
<code>-Gy</code>	Enables reordering of functions at link time.	ON	*
<code>-H n</code>	Limits the length of external symbol names to <code>n</code> characters.	OFF	*
<code>-Idirectory</code>	Specifies an additional directory to search for include files.	OFF	page 3-4
<code>-J</code>	Makes the default <code>char</code> type unsigned.	OFF	*
<code>-LD</code>	Produces a DLL.	OFF	*
<code>-LDd</code>	Creates a debug version of DLL.	OFF	*
<code>-link</code>	Passes options to the linker.	OFF	page 3-5
<code>-MD</code>	Compiles and links with the dynamic, multi-thread C runtime library.	OFF	*
<code>-MDd</code>	Compiles and links with the dynamic, multi-thread, debug version of the C runtime library.	OFF	*
<code>-ML</code>	Compiles and links with the static, single-thread C runtime library.	ON	*
<code>-MLd</code>	Compiles and links with the static, single-thread, debug version of the C runtime library.	OFF	*

* The options marked with the asterisk in the "Reference" column produce the same results for the Intel C/C++ Compiler as for the Microsoft Visual C++ Compiler. If you need more information on these options, see the documentation that accompanies Microsoft Visual C++.

Table 2-2 Summary of Command-line Options (Continued)

Option	Description	Default	Reference
<code>-MT</code>	Compiles and links with the static, multi-thread C runtime library.	OFF	*
<code>-MTd</code>	Compiles and links with the static, multi-thread, debug version of the C runtime library.	OFF	*
<code>-O1</code>	Optimize for speed, but disable some optimizations that increase code size for small speed benefit. The <code>-O1</code> option has the same effect as specifying the options: <code>-Og</code> , <code>-Oi</code> , <code>-Os</code> , <code>-Oy</code> , <code>-Ob1</code> , <code>-Gf</code> , <code>-Gs</code> , and <code>-Gy</code> .	OFF	page 4-1
<code>-O2</code>	Optimizes for speed. The <code>-O2</code> option has the same effect as specifying the following options: <code>-Og</code> , <code>-Oi</code> , <code>-Ot</code> , <code>-Oy</code> , <code>-Ob1</code> , <code>-Gf</code> , <code>-Gs</code> , and <code>-Gy</code> .	ON	page 4-1
<code>-Obn</code>	Controls the compiler's inline expansion. The amount of inline expansion performed varies with the value of <code><n></code> as follows. <code>0</code> Disables inlining. <code>1</code> Enables inlining of functions declared with the <code>_inline</code> keyword. Also enables inlining according to the C++ language. <code>2</code> Enables inlining of any function. However, the compiler decides which functions to inline. Enables interprocedural optimizations and has the same effect as <code>-Qip</code> .	OFF	page 5-4
<code>-Od</code>	Disables optimizations.	OFF	page 4-1
<code>-Og</code>	Enables global optimizations.	ON	*
<code>-Oi</code>	Enables inline expansion of standard library functions.	ON	page 4-6

* The options marked with the asterisk in the "Reference" column produce the same results for the Intel C/C++ Compiler as for the Microsoft Visual C++ Compiler. If you need more information on these options, see the documentation that accompanies Microsoft Visual C++.

Table 2-2 Summary of Command-line Options (Continued)

Option	Description	Default	Reference
<code>-Oi-</code>	Disables inline expansion of standard library functions.	OFF	page 4-6
<code>-Op</code>	Favors conformance to the ANSI C and IEEE 754 standards for floating-point arithmetic. Behavior for NaN comparisons does not conform.	OFF	page 4-7
<code>-Op-</code>	Favors optimization over conformance to the ANSI C and IEEE 754 standards for floating-point arithmetic.	ON	page 4-7
<code>-Os</code>	Enables most speed optimizations, but disable optimizations that increase code size for a small speed benefit.	OFF	*
<code>-Ot</code>	Enables all speed optimizations.	ON	*
<code>-Ox</code>	Same as the -O2 option without the -Gf and -Gy options.	OFF	*
<code>-Oy</code>	Enables the use of the ebp register in optimizations.	ON	page 6-6
<code>-Oy-</code>	Disables the use of the ebp register in optimizations. Instead, it is used as the frame pointer.	OFF	page 6-6
<code>-P</code>	Stops the compilation process after C or C++ source files have been preprocessed and writes the results to files named according to the compiler's default file-naming conventions.	OFF	page 7-2
<code>-QA-</code>	Causes all predefined macros (other than those beginning with <code>__</code>) and all assertions to be inactive. (Same as <code>-u.</code>)	OFF	page 7-3

* The options marked with the asterisk in the "Reference" column produce the same results for the Intel C/C++ Compiler as for the Microsoft Visual C++ Compiler. If you need more information on these options, see the documentation that accompanies Microsoft Visual C++.

Table 2-2 Summary of Command-line Options (Continued)

Option	Description	Default	Reference
<code>-QAname[(values)]</code>	Associates a symbol name with the specified sequence of values. Equivalent to an <code>#assert</code> preprocessing directive.	OFF	page 7-3
<code>-Qax[extensions]</code>	Generates specialized and generic code for processor-specific extensions.	OFF	page 4-4
<code>-QH</code>	Produces a list of include file dependencies, one per line, to <code>stdout</code> .	OFF	page 7-5
<code>-QIOf</code>	Avoids the incorrect decoding of certain <code>Of</code> instructions for code targeted at older processors.	OFF	page 11-4
<code>-QIfdiv</code>	Enables a software patch for the floating-point division flaw that exists in some steppings of the Pentium processor.	OFF	page 11-3
<code>-QIfdiv-</code>	Disables the software patch for the floating-point division flaw that exists in some steppings of the Pentium processor.	ON	page 11-3
<code>-Qip</code>	Enables interprocedural optimizations.	OFF	page 5-2
<code>-Qip_no_inlining</code>	Disables inlining that would result from the <code>-Qip</code> or <code>-Ob2</code> interprocedural optimizations, but has no effect on other interprocedural optimizations. This option has no effect on user-directed inlining (<code>-Ob1</code>).	OFF	page 5-4
<code>-Qipo</code>	Enables interprocedural optimization across files.	OFF	page 5-2
<code>-Qkscalar</code>	Performs all 32-bit floating point arithmetic using the Streaming SIMD Extensions instead of the default x87 instructions.	OFF	page 14-3

Table 2-2 Summary of Command-line Options (Continued)

Option	Description	Default	Reference
<code>-Qlocation,tool,path</code>	Specifies an alternate version of a <code>tool</code> specified by <code>path</code> .	OFF	page 3-4
<code>-Qlong_double</code>	Changes the default size of the <code>long double</code> data type from 64 to 80 bits.	OFF	page 4-10
<code>-QM</code>	Generates makefile dependency lines for each source file, based on the <code>#include</code> lines found in the source file.	OFF	page 7-6
<code>-Qnobss_init</code>	Places variables that are initialized with zeroes in the DATA section.	OFF	page 12-2
<code>-Qoption,tool,list</code>	Passes an argument list to another program in the compilation sequence, such as the assembler or linker.	OFF	page 3-4
<code>-Qpc nn</code>	Enables floating-point significand precision control. The value of <code>nn</code> is used to round the significand to the correct number of bits. The value of <code>nn</code> must be either 32, 64, or 80.	<code>nn=64</code>	page 4-8
<code>-Qpf[options]</code>	Improves cache usage with prefetching.	OFF	page 4-5
<code>-Qprec</code>	Improves floating-point precision but does not provide the speed benefit of <code>-Op</code> .	OFF	page 4-8
<code>-Qprec_div</code>	Disables the floating point division-to-multiplication optimization.	OFF	page 4-8
<code>-Qprof_dir dirname</code>	Specifies the directory to hold profile information.	OFF	page 5-5
<code>-Qprof_gen</code>	Instruments the program to collect basic block execution counts.	OFF	page 5-5
<code>-Qprof_genx</code>	Generates an instrumented object file and also creates a new static profile information file (<code>.spi</code>).	OFF	page 5-5
<code>-Qprof_use</code>	Uses dynamic feedback information.	OFF	page 5-5

Table 2-2 Summary of Command-line Options (Continued)

Option	Description	Default	Reference
<code>-Qrcd</code>	Disables changing of the FPU rounding control.	OFF	page 4-10
<code>-Qrestrict</code>	Enables pointer disambiguation with the <code>restrict</code> qualifier.	OFF	page 14-3
<code>-Qsfsalign-</code>	Disables stack alignment for all functions.	OFF	page 13-44
<code>-Qsfsalign16</code>	Aligns stack for functions with 16 byte variables.	OFF	page 13-44
<code>-Qsfsalign8</code>	Aligns stack for functions with 8 or 16 byte variables.	ON	page 13-44
<code>-Qunrolln</code>	Specifies the maximum number of times to unroll a loop.	<code>n = 0</code>	page 4-6
<code>-Quse_asm</code>	Compiles source files to assembly files and calls MASM to generate object files automatically.	OFF	page 6-4
<code>-Qvc4</code>	Enables compatibility with Visual C++ 4.x.	OFF	page 9-2
<code>-Qvc5</code>	Enables compatibility with Visual C++ 5.0.	ON	page 9-2
<code>-Qvc6</code>	Enables compatibility with Visual C++ 6.0.	ON	page 9-2
<code>-Qvec</code>	Enables the vectorizer.	OFF	page 14-2
<code>-Qvec_alignment</code>	Controls the default alignment of vectorizable data.	OFF	page 14-2
<code>-Qvec_emms[-]</code>	Controls whether the compiler automatically inserts an EMMS instruction after vectorization of loops.	OFF	page 14-3
<code>-Qvec_no_alias[-]</code>	Assumes that no aliasing can occur between objects with different names.	OFF	page 14-4
<code>-Qvec_no_arg_alias[-]</code>	Assumes on entry that procedure arguments are not aliased.	OFF	page 14-4
<code>-Qvec_verbose</code>	Controls the diagnostic levels produced by the vectorizer.	OFF	page 14-3

Table 2-2 Summary of Command-line Options (Continued)

Option	Description	Default	Reference
<code>-Qwdtag</code>	Disables the soft diagnostic that corresponds to <i>tag</i> .	OFF	page 10-5
<code>-Qwe[tag]</code>	Overrides the severity of the soft diagnostic corresponding to <i>tag</i> and makes it an error.	OFF	page 10-5
<code>-Qwnnum</code>	Limits the number of errors displayed prior to aborting compilation to <i>num</i> .	100	page 10-6
<code>-Qwr[tag]</code>	Overrides the severity of the soft diagnostic corresponding to <i>tag</i> and makes it a remark.	OFF	page 10-5
<code>-Qww[tag]</code>	Overrides the severity of the soft diagnostic corresponding to <i>tag</i> and makes it a warning.	OFF	page 10-5
<code>-Qx[extensions]</code>	Generates specialized code for processor-specific extensions.	OFF	page 4-4
<code>-S</code>	Stops the compilation process after an assembly source has been generated. The compiler writes assembly code source for each C or C++ source file or preprocessed source file to a file using the output file's naming conventions.	OFF	page 6-2
<code>-TC</code>	Compiles all source or unrecognized file types as C source files.	OFF	*
<code>-Tc file</code>	Treats <i>file</i> as a C source file.	OFF	*
<code>-TP</code>	Compiles all source or unrecognized file types as C++ source files.	OFF	*
<code>-Tp file</code>	Treats <i>file</i> as a C++ source file.	OFF	*
<code>-u</code>	Disables all predefined macros (other than those starting with <code>__</code>) and assertions. (Same as <code>-QA-</code> .)	OFF	*

* The options marked with the asterisk in the "Reference" column produce the same results for the Intel C/C++ Compiler as for the Microsoft Visual C++ Compiler. If you need more information on these options, see the documentation that accompanies Microsoft Visual C++.

Table 2-2 Summary of Command-line Options (Continued)

Option	Description	Default	Reference
<code>-Uname</code>	Suppresses any definition of a macro name; equivalent to a <code>#undef</code> preprocessing directive.	OFF	page 7-3
<code>-V text</code>	Sets version string to <code>text</code> .	OFF	*
<code>-vd0</code>	Suppresses C++ hidden <code>vtordisp</code> constructor/destructor displacement member.	OFF	*
<code>-vd1</code>	Generates C++ hidden <code>vtordisp</code> constructor/destructor displacement member.	ON	*
<code>-vmb</code>	Selects the smallest representation for pointers to members. Use this option if you define each class before you declare a pointer to a member of the class.	ON	*
<code>-vmg</code>	Selects the general representation for pointers to members. Use this option if you declare a pointer to a member before you define the corresponding class.	OFF	*
<code>-vmm</code>	Enables pointers to class members with single or multiple inheritance with <code>-vmg</code> .	OFF	*
<code>-vms</code>	Enables pointers to members of single-inheritance classes with <code>-vmg</code> .	OFF	*
<code>-vmv</code>	Enables pointers class members of any inheritance type with <code>-vmg</code> .	OFF	*
<code>-W0</code>	Specifies to only display error messages (no warnings or remarks).	OFF	page 10-4
<code>-W1, -W2, -W3</code>	Specifies that error and warning messages are displayed.	ON	page 10-5

* The options marked with the asterisk in the "Reference" column produce the same results for the Intel C/C++ Compiler as for the Microsoft Visual C++ Compiler. If you need more information on these options, see the documentation that accompanies Microsoft Visual C++.

Table 2-2 Summary of Command-line Options (Continued)

Option	Description	Default	Reference
<code>-W4</code>	Specifies that error, warning, and remark messages are displayed.	OFF	page 10-6
<code>-X</code>	Removes the standard directories from the list of directories to be searched for include files.	OFF	page 3-3
<code>-Za</code>	Enforces strict conformance to the ANSI standard for C.	OFF	page 8-2
<code>-Zd</code>	Produces only line number information in the object file (for debugging).	OFF	*
<code>-Ze</code>	Accepts extended C language.	ON	page 8-2
<code>-Zi</code>	Generates symbolic debugging information in the object code for use by source-level debuggers.	OFF	page 6-6
<code>-Zl</code>	Disables embedding default libraries in object files.	OFF	*
<code>-Zp[number]</code>	Specifies the strictest alignment constraint for structure and union types as one of the following: 1, 2, 4, 8, or 16 bytes.	8	page 12-1
<code>-Zs</code>	Checks the syntax of a program and stops the compilation process after the C or C++ source files and preprocessed source files have been parsed. Generates no code and produces no output files. Warnings and messages appear on <code>stderr</code> .	OFF	page 6-2

* The options marked with the asterisk in the "Reference" column produce the same results for the Intel C/C++ Compiler as for the Microsoft Visual C++ Compiler. If you need more information on these options, see the documentation that accompanies Microsoft Visual C++.

Default Behavior of the Compiler

You need not enter any options when you invoke the Intel C/C++ compiler. If you do not specify any options, the compiler uses default settings. The compiler driver performs the following actions by default:

- Produces executable output with the filename of the first source or object file on the command line with a `.exe` suffix.
- Optimizes for the Pentium Pro or Pentium II processors with `-G6`.
- Searches for include files using the `INCLUDE` variable.
- Searches for library files in directories specified by the `LIB` variable, if they are not found in the current directory.
- Sets 16 bytes as the strictest alignment constraint for structures.
- Displays error and warning messages.
- Uses ANSI with extensions (`-Ze`) for C source files.
- Performs standard optimizations using the default `-O2` option, as described in [“Optimization Choices” in Chapter 4](#).

If the compiler does not recognize a command-line option, that option is ignored and a warning is displayed. See, [Chapter 10, Diagnostic Information](#) for detailed descriptions about system messages.

Changing the Compilation Environment

3

This chapter describes how to customize the environment you use during compilation. Through customizing you can specify the following:

- names and locations of compilation tools
- options to include during each compilation
- options and files to use for individual projects
- directories in which to search for include files and libraries

Environment Variables

To customize your environment, you can specify paths where the compiler can search for special files such as libraries or include files. On Win32* systems, you can specify these variables in the Environment tab of the System Properties window. You can access this window by clicking on the System icon in the Control Panel. The following variables are relevant to your compilation environment:

- **LIB** specifies the directory path for the math libraries
Also, the compiler calls **LINK.exe**, the Microsoft linker, to produce an executable file from the object files. This linker searches the path specified in the **LIB** environment variable to find the libraries.
- **PATH** specifies the directory path for the compiler executable files
- **INCLUDE** specifies the directory path for the for include files.
- **TMP** specifies the directory in which to store temporary files. If the directory specified by **TMP** does not exist, the compiler places the temporary files in the current directory.

Configuration Files

A user configuration file named `icl.cfg` can be added to the directory where `icl.exe` is installed. In this file, you can specify common options that you want to include in each compilation. Options in the configuration file are processed in the order they appear, followed by the command-line options that you specify when you invoke the compiler.

You can decrease the time you spend entering command-line options and ensure consistency by using the configuration file to automate some command-line entries. However, you should be aware that options placed in the configuration file will be included each time you run the compiler. If you have varying option requirements for different projects, see [“Response Files”](#) later in this chapter.

You can insert any valid command-line option into this file. Use the pound “#” character to treat the rest of the line as a comment.

[Example 3-1](#) illustrates a sample `icl.cfg` file.

Example 3-1 Sample icl.cfg File

```
## Sample icl.cfg file.
##
## Define preprocessor macro MY_PROJECT.
-DMY_PROJECT
##
## Additional directories to be searched for include
## files, before the default.
-Ic:\project\include
##
## Use the static, multi-threaded C run-time library.
-MT
```

Response Files

You can use response files to specify options used during particular compilations and save this information in individual files. Response files are invoked as an option on the command line. Options in a response file are inserted in the command line at the point where the response file is invoked.

You can decrease the time you spend entering command line options and ensure consistency by using response files to automate command line entries. Further, you can use individual response files to maintain options for specific projects and avoid editing the configuration file when changing from one project to the next.

Any number of options or filenames can be placed on a line in the response file. You can also use several response files in the same command line. You can insert any valid command-line option into this file. Use the pound “#” character to treat the rest of the line as a comment.

The syntax for using response files is as follows:

```
prompt> icl @response_file filenames
```

Note that an “at” sign (@) must precede the name of the response file on the command line.

Include Files

By default, the compiler searches for the standard include files in the directories specified in the `INCLUDE` environment variable. You can also indicate the location of include files in the configuration file, `icl.cfg`, using the options described in [“Specifying an Include Directory \(-I\)”](#).

Specifying an Include Directory (-I)

Use the `-Idirectory` option to specify an additional directory in which to search for include files. Included files are brought into the program with a `#include` preprocessor directive. The compiler searches directories for include files in the following order:

- directory of the source file that contains the `include`
- directories specified by the `-I` option
- directories specified in the `INCLUDE` environment variable

Removing Include Directories (-X)

Use the `-X` option to prevent the compiler from searching the default path specified by the `INCLUDE` environment variable.

You can use the `-X` option with the `-I` option to prevent the compiler from searching the default path for include files and direct it to use an alternate path. For example, to direct the compiler to search the path `\alt\include` instead of the default path, do the following:

```
prompt> icl -X -I \alt\include newmain.cpp
```

How To Specify Alternate Tools and Paths

The Intel compiler lets you go outside the default paths and tools to specify alternate tools for preprocessing, compilation, assembly, and linking. Further, you can invoke options specific to your alternate tools on the command line. This functionality is provided by `-Qlocation` and `-Qoption`.

Specifying an Alternate Component (`-Qlocation,tool,path`)

Use `-Qlocation` to specify an alternate path for a tool. This option accepts two arguments using the following syntax:

```
prompt> -Qlocation,tool,path
```

`tool` Indicates one of the following to designate the compilation tool to receive one or more of these arguments:

<code>cpp</code>	Specifies the compiler front-end preprocessor.
<code>c</code>	Specifies the C/C++ compiler.
<code>asm</code>	Specifies the assembler.
<code>link</code>	Specifies the linker.

Passing Options to Other Programs (`-Qoption, tool, optlist`)

Use `-Qoption` to pass an option specified by `optlist` to a `tool`, where `optlist` is a comma-separated list of options. The syntax for this command is the following;

```
Qoption,tool,optlist
```

<i>tool</i>	Indicates one of the following that designate the compilation tool to receive one or more of these arguments:
<i>cpp</i>	Specifies the compiler front-end preprocessor.
<i>c</i>	Specifies the C/C++ compiler.
<i>asm</i>	Specifies the assembler.
<i>link</i>	Specifies the linker.
<i>optlist</i>	Indicates one or more valid argument strings for the designated program. If the argument is a command-line option, you must include the hyphen. If the argument contains a space or tab character, you must enclose the entire argument in quotation characters (" "). You must separate multiple arguments with commas.

The following example directs the linker to create a memory map when the compiler produces the executable file from the source.

```
prompt> icl -Qoption,link,-map:proto.map proto.cpp
```

The `-Qoption,link` option in the preceding example is passing the `-map` option to the linker. This is an explicit way to pass arguments to other tools in the compilation process.

To get the same results implicitly you could also use the following:

```
prompt> icl -Fmproto.map proto.cpp
```

Passing Options to the Linker (-link)

To pass options specifically to the linker, use the `-link` option. For example, to compile the file `a.cpp` and instruct the linker to link `a.obj` with `libfoo.lib`, creating `a.exe`, type the following at the prompt:

```
prompt> icl a.cpp -link libfoo.lib
```



NOTE. *The compiler passes everything following `-link` to the linker only. Therefore, all other compiler options must precede `-link`.*

Optimizations

4

This chapter describes ways to improve the performance of your application by using the optimization options on programs. You can find related information on optimizations in [Chapter 5, Interprocedural and Profile Guided Optimizations](#) and [Appendix B, “Experimental Performance Tuning.”](#) which describes some experimental ways to tune your code.

Optimization Choices

To specify the kind of optimization you want for your program, use the following command-line options:

- | | |
|------------------|--|
| <code>-Od</code> | Disables optimizations. |
| <code>-O1</code> | Enables options <code>-Og</code> , <code>-Oi-</code> , <code>-Os</code> , <code>-Oy</code> , <code>-Ob1</code> , <code>-Gf</code> , <code>-Gs</code> , and <code>-Gy</code> . However, <code>-O1</code> disables a few options that increase code size for a small speed benefit such as inline expansion of library functions. This option is not as aggressive as the Microsoft Visual C++ Compiler version of the <code>-O1</code> option. In most cases, <code>-O2</code> is recommended over <code>-O1</code> because the <code>-O2</code> option enables inline expansion, which helps programs that have many function calls. |
| <code>-O2</code> | Enables options <code>-Og</code> , <code>-Oi</code> , <code>-Ot</code> , <code>-Oy</code> , <code>-Ob1</code> , <code>-Gf</code> , <code>-Gs</code> , and <code>-Gy</code> . Confines optimizations to the procedural level. See Table 4-1 for a comprehensive list of optimizations. The <code>-O2</code> option is on by default. |

Restricting Optimization with -Od

The `-Od` option disables optimizations. [Table 4-1](#) shows the optimizations that the compiler applies when you invoke the optimization in the **Option** column. The only optimization not disabled by `-Od` is the one that has “any” in the **Option** column.

Table 4-1 Optimization Summary

Optimization	Affected Aspect of Program	Option
constant propagation	constants and expression evaluation	-O1 / -O2
copy propagation	constants and expression evaluation	-O1 / -O2
dead-code elimination	instruction sequencing	-O1 / -O2
global register allocation	register use	-O1 / -O2
instruction scheduling	instruction reordering	-O1 / -O2
loop unrolling	instruction sequencing	-O1 / -O2
loop-invariant code movement	instruction sequencing	-O1 / -O2
optimized code selection	instruction selection/ addressing modes	any
partial redundancy elimination	constants and expression evaluation	-O1 / -O2
strength reduction/induction variable simplification	instruction selection/sequencing constants and expression evaluation	-O1 / -O2
variable renaming	register use	-O1 / -O2

The Effect of -Od on -Oy and Debugging

When you specify `-Od`, you automatically disable `-Oy`, which is otherwise enabled by default or when you specify `-O1` or `-O2`. The `-Oy` option allows the compiler to use the `ebp` register as a general purpose register in optimizations. However, most debuggers expect `ebp` to be used as a stack frame pointer, and cannot produce a stack backtrace unless this is so. The `-Oy-` option instructs the compiler to generate code that maintains and uses `ebp` as a stack frame pointer, without turning off other `-O1` or `-O2` optimizations. As a result, the debugger can still produce a stack backtrace. However, using `-Oy-` reduces the number of available general-purpose registers by one, and can result in slightly less efficient code.

Targeting a Processor (-Gn)

Use `-Gn` to target an application to run on a specific processor for maximum performance. Regardless of which of the `-Gn` suboptions you choose, the resulting binary will run on any Intel Architecture 32-bit (IA-32) processor. The following are the suboptions for `-Gn`:

- `G4` Choose `G4` when the binary application runs primarily on the Intel486 processor.
- `G5` Choose `G5` when the binary application runs primarily on the Pentium processor.
- `G6` Choose `G6` when the binary application runs primarily on the Pentium Pro or Pentium II processors. The `G6` switch is ON by default.

For example, the following commands both compile the source program `prog.cpp` and optimize for the Pentium Pro and Pentium II processors:

```
prompt> icl prog.cpp
prompt> icl -G6 prog.cpp
```

Automatic Processor Dispatch Support (-Qx[extensions] and -Qax[extensions])

The `-Qx[extensions]` and `-Qax[extensions]` options provide support to generate code that is specific to processor-instruction extensions.

`-Qx[extensions]` generates specialized code to run exclusively on the processors indicated by the extension.

`-Qax[extensions]` generates code specialized to the specified extensions, but also generates generic IA-32 code. The generic code is usually slower.

You can specify the same extensions for either option as listed in [Table 4-2](#).

Table 4-2 Processor Dispatch Extension Options

Extension	Processors and Features Provided by the Extension
i	Pentium® Pro processor and Pentium II processors, which use the CMOV and FCMOV instructions
M	Pentium Pro processors with MMX™ technology instructions
K	Pentium III processor with the Streaming SIMD Extensions, which include the i and M instruction sets as well

Exclusive Specialized Code with **-Qx[extensions]**

The **-Qx[extensions]** option specifies the minimum extension set required to exist on processors on which you execute your program. The resulting code can contain unconditional use of the specified processor extensions. When you use **-Qx[extensions]**, the code generated by the compiler cannot execute on IA-32 processors that lack the specified extensions. The following example compiles the program contained in **myprog.cpp**, using the **i** extension:

```
prompt> icl -O2 -G6 -Qxi myprog.cpp
```

The resulting program, **myprog.exe**, might not execute on a Pentium processor, but will execute on Pentium Pro and Pentium II processors.



CAUTION. A program compiled with **-Qx[extensions]** can fail with an *Illegal instruction exception*, or other unexpected behavior, if it is executed on a processor that lacks the specified extensions.

Specialized and Generic Code with **-Qax[extensions]**

Use **-Qax[extensions]** to enable automatic code specialization for a specific set of processors extensions. The compiler generates code to detect, at runtime, the extensions supported by the current processor. Specialized code is executed if the applicable extensions are supported by the processor. The compiler also generates an equivalent, but likely slower, version of code that is generic. Thus, programs compiled with **-Qax[extensions]** execute on any IA-32 processor, subject to any restrictions imposed if you also compile with the **-Qx[extensions]**.

The following command compiles the program contained in `myprog.cpp`, and allows the compiler to generate specialized code to exploit the architectural features of the Pentium Pro processor such as the `CMOV` instruction and the MMX™ technology instructions:

```
prompt> icl -O2 -G6 -QaxiM myprog.cpp
```

The resulting program, `myprog.exe`, executes on all IA-32 processors.



CAUTION. When you use `-Qax[extensions]` in conjunction with `-Qx[extensions]`, the extensions specified by `-Qx[extensions]` can be used unconditionally by the compiler, and the resulting program will require the processor extensions to execute properly.

Prefetching with `-Qpf[options]`

Use `-Qpf` to automatically insert prefetching on a Pentium III processor. This option enables three suboptions described in the list that follows, each of which seeks to improve cache behavior. The following example invokes `-Qpf` as one option with all its functionality:

```
prompt> icl -Qpf[options] a.cpp
```

Invoke the following suboptions to provide only some of the functionality:

- | | |
|------------------------|--|
| <code>-Qpf_loop</code> | Automatically inserts prefetch instructions for array references that have as their subscript a simple linear function of the index variable. In addition it will improve this prefetching efficiency by unrolling loops so that 1 prefetch is executed for each iteration of the resulting loop or the size of the loop body is sufficiently large to cover the possible latency of a prefetch. The number of iterations unrolled depends on the size of the data elements, the stride of the loop and the size of the loop body. |
| <code>-Qpf_call</code> | When you make a call to a function with a pointer to a <code>struct</code> as an argument, <code>-Qpf</code> assumes that the struct will be used within the function, and so it generates a prefetch of that <code>struct</code> . |

-Qpf_sstore Enables the generation of streaming store instructions. Streaming store instructions reduce the chance of cache pollution, and can improve usage of the memory bus.

Loop Unrolling with -Qunrolln

Use **-Qunrolln** to specify the maximum number of times you want to unroll a loop. The following example unrolls a loop at most four times:

```
prompt> icl -Qunroll4 a.cpp
```

To disable loop unrolling, specify *n* as 0. The following example disables loop unrolling:

```
prompt> icl -Qunroll0 a.cpp
```

Inline Expansion of Library Functions (-Oi, -Oi-)

By default, the compiler inlines a number of standard C, C++, and math library functions. This usually results in faster execution of your program.

Sometimes inline expansion of library functions can cause unexpected results. The inlined library functions do not set the **errno** variable. So in code that relies upon the setting of the **errno** variable, you should use the **-Oi-** option. Also, if one of your functions has the same name as one of the compiler's supplied library functions, the compiler assumes that it is one of the latter and replaces the call with the inlined version. Consequently, if the program defines a function with the same name as one of the known library routines, you must use the **-Oi-** option to ensure that the program's function is the one used.

For example, use **-Oi-** to ensure the disabling of inline expansion with the following options:

- IEEE 754 conformance (**-Op**)
- ANSI standard conformance (**-Za**)
- restrict optimization (**-Od**)
- debugging information (**-Zi**)

Your results can vary slightly using the preceding optimizations.



NOTE. *Automatic inline expansion of library functions is not related to the inline expansion that the compiler does during interprocedural optimizations. For example, the following command compiles the program `sum.cpp` without expanding the library functions:*

```
prompt> icl -Qip -Oi- sum.cpp
```

Floating-point Arithmetic Precision (-Op, -Op-, -Qprec, -Qprec_div, -Qpc, -Qlong_double)

The options described in this section all provide optimizations with varying degrees of precision in floating-point arithmetic.

When to Use -Op

The `-Op` option restricts optimization to maintain declared precision and to ensure that floating-point arithmetic conforms more closely to the ANSI and IEEE standards.

For most programs, specifying this option adversely affects performance. If you are not sure whether your application needs this option, try compiling and running your program both with and without it to evaluate the effects on performance versus precision.

Specifying this option has the following effects on program compilation:

- User variables declared as floating-point types are not assigned to registers.
- Whenever an expression is spilled, it is spilled as 80 bits (extended precision), not 64 bits (double precision).
- Floating-point arithmetic comparisons conform to IEEE 754 except for NaN behavior.
- The exact operations specified in the code are performed. For example, division is never changed to multiplication by the reciprocal.

- The compiler performs floating-point operations in the order specified without reassociation.
- The compiler does not perform the constant-folding optimization on floating-point values. Constant folding also eliminates any multiplication by 1, division by 1, and addition or subtraction of 0. For example, code that adds 0.0 to a number is executed exactly as written. Compile-time floating-point arithmetic is not performed to ensure that floating-point exceptions are also maintained.
- Floating-point operations conform to ANSI C. When assignments to type `float` and `double` are made, the precision is rounded from 80 bits (extended) down to 32 bits (`float`) or 64 bits (`double`). When you do not specify `-Op`, the extra bits of precision are not always rounded before the variable is reused.
- The `-Oi` option, which disables inline functions expansion, is used.



NOTE. The `-Oi` and `-Op` options are active by default when you choose the `-Za` (strict ANSI C conformance) option.

When to Use `-Qprec`

Use the `-Qprec` option to improve floating-point precision. `-Qprec` disables fewer optimizations and has less impact on performance than `-Op`.

When to Use `-Qprec_div`

Use `-Qprec_div` to disable the floating point division-to-multiplication optimization. The Intel C/C++ Compiler can change floating-point division computations into multiplication by the reciprocal of the denominator. This change can alter the results of floating point division computations slightly.

When to Use `-Qpc`

Use the `-Qpc nn` option to enable floating-point significand precision control. Some floating-point algorithms are sensitive to the accuracy of the significand or fractional part of the floating-point value. For example,

iterative operations like division and finding the square root can run faster if you lower the precision with the `-Qpc nn` option. Set `nn` to one of the following values to round the significand to the indicated number of bits:

`-Qpc 32` 24 bits (single precision)
`-Qpc 64` 53 bits (double precision)
`-Qpc 80` 64 bits (extended precision)

The default value for `nn` is 64, indicating double precision.

This option allows full optimization. Using this option does not have the negative performance impact of using the `-Op` option because only the fractional part of the floating-point value is affected. The range of the exponent is not affected.

The `-Qpc nn` option causes the compiler to change the floating point precision control when the `main()` function is compiled. The program that uses `-Qpc nn` must use `main()` as its entry point, and the file containing `main()` must be compiled with `-Qpc nn`.

Alternatives to the `-Qpc` Option

The `-Qpc nn` option performs its intended function if your program has a `main()` routine and only if you use the Intel compiler to compile the file that contains `main()`. Otherwise, this option is not useful.

An alternative method to achieve the effect of `-Qpc nn`, which does not require the use of the Intel compiler to compile `main()`, is through the use of the Microsoft C library routine `_controlfp()`. The `_controlfp()` function enables floating-point precision modification. To use this routine, you must include the Microsoft header file `float.h`. [Table 4-3](#) summarizes how to achieve the effects of `-Qpc nn` using `_controlfp()`.

Table 4-3 `-Qpc` Equivalent Calls

Option	Equivalent Call	Extra Options
<code>-Qpc 32</code>	<code>_controlfp(_PC_24, _MCW_PC);</code>	none
<code>-Qpc 64</code>	<code>_controlfp(_PC_53, _MCW_PC);</code>	none
<code>-Qpc 80</code>	<code>_controlfp(_PC_64, _MCW_PC);</code>	none



NOTE. For better performance, it is preferable to insert the call to `_controlfp()` within an initialization routine in your program.

[Figure 4-1](#), shows how to use `_controlfp()` in a program. See the Microsoft On-line Help for more information on `_controlfp()`.

Figure 4-1 Example of the `_controlfp()` Function:

```
// init.c - contains initialization code for my application
#include <float.h>
void init_program(void){
// Set FPU precision control to use 24-bit significand
_controlfp( _PC_24, _MCW_PC );

// other initialization code follows...
}
```

When to Use `-Qlong_double`

Use the `-Qlong_double` to change the size of the `long double` type to 80-bits. For compatibility with the Microsoft compiler, the Intel compiler's default `long double` type is 64 bits in size, the same as the `double` type. This option introduces a number of incompatibilities with other files compiled without this option and with calls to library routines. Therefore, Intel recommends that the use of `long double` variables be local to a single file when you compile with this option.

Rounding Control Option (`-Qrcd`)

The Intel compiler uses the `-Qrcd` option to improve the performance of code that requires floating point calculations. The optimization is obtained by controlling the change of the rounding mode.

The system default floating point rounding mode is round-to-nearest. This means that values are rounded during floating point calculations. However, the C language requires floating point values to be truncated when a

conversion to an integer is involved. To do this, the compiler must change the rounding mode to truncation before each floating point conversion and change it back afterwards.

The `-Qrcd` option disables the change to truncation of the rounding mode in floating point-to-integer conversions. This means that all floating point calculations must use the default round-to-nearest, including floating point-to-integer conversions. This option has no effect on floating point calculations, but conversions to integer will not conform to C semantics.

Interprocedural and Profile Guided Optimizations

5

This chapter describes two methods to improve the performance of your code based on its unique profile and procedural dependencies.

Interprocedural Optimization (IPO)—Use the `-Qip` option to analyze your code and apply optimizations between procedures within each source file. Use multifile IPO with `-Qipo` to enable the optimizations between procedures in separate source files.

Profile Guided Optimization (PGO)—Creates an instrumented program from your source code and special code from the compiler. Each time this instrumented code is executed, the compiler generates a dynamic information file. When you compile a second time, the dynamic information files are merged into a summary file. Using the profile information in this file, the compiler attempts to optimize the execution of the most heavily travelled paths in the program.

Unlike other optimizations such as those strictly for size or speed, the results of IPO and PGO vary. This is due to each program having a different profile and different opportunities for optimizations. The guidelines provided help you determine if you can benefit by using IPO and PGO. But, only you can make this evaluation by understanding the principles of the optimizations and the unique aspects of your source code.



CAUTION. The `-Qip` and `-Qipo` options can in some cases significantly increase compile time and code size.

Interprocedural Optimization (IPO) with -Qip and -Qipo

Use `-Qip` to enable interprocedural optimizations, which allow the compiler to analyze your code to determine where you can benefit from the optimizations listed in [Table 5-1](#).

Table 5-1 IPO Optimization Summary

Optimization	Affected Aspect of Program
inline function expansion	calls, jumps, branches, and loops
interprocedural constant propagation	arguments, global variables, and return values
passing arguments in registers	calls, register usage
loop-invariant code motion	further optimizations, loop invariant code
dead code elimination	code size
propagation of function characteristics	call deletion and call movement

Inline function expansion is one of the main optimizations performed by the interprocedural optimizer. For function calls that the compiler believes are frequently executed, the compiler might decide to replace the instructions of the call with code for the function itself.

With `-Qip`, the compiler performs inline function expansion for calls within procedures defined within the current source file. However, when you use `-Qipo` to specify multifile IPO, the compiler performs inline function expansion for calls between procedures defined in separate files.

Multifile IPO (-Qipo)

Multifile IPO obtains potential optimization information from individual program modules of a multifile program. Using the information, the compiler performs optimizations across modules. The result of `-Qipo` is a set of information files with a `.il` suffix.

Building a program is divided into two phases: compilation and linkage. Multifile IPO performs different work depending on whether the compilation or linkage is being performed.

1. *Compilation Phase*—As each source file is compiled, multifile IPO creates an information file. This file has the same name as the object file but with a `.il` suffix. The information file contains an intermediate representation of the source code as well as summary information used for optimization. The compiler also produces a `.obj` file when multifile IPO is enabled.
2. *Linkage Phase*—When you specify `-Qipo`, the compiler is invoked a final time before the linker. The compiler performs multifile IPO across all modules for which a corresponding up-to-date `.il` file exists.

Creating a Multifile IPO Executable

To following two steps show how to enable multifile IPO:

1. Compile your modules with `-Qipo`; as follows:

```
prompt> icl -Qipo -c a.cpp b.cpp c.cpp
```

Use `-c` to stop compilation after generating the `.il` and `.obj` files.

Resulting information files: `a.il`, `b.il`, and `c.il`

Resulting object files: `a.obj`, `b.obj`, and `c.obj`



NOTE. The path and name of the `.il` file must match the object file.

2. With preceding results, you can now optimize interprocedurally:

```
prompt> icl -Fenu_ip_ofile -Qipo a.obj b.obj c.obj
```

The `-Fe` option stores the executable in `nu_ipo_file.exe`.

Multifile IPO is applied only to modules with up-to-date `.il` files, otherwise the object file passes to link stage.

For efficiency, combine steps 1 and 2:

```
prompt> icl -Qipo -Fenu_ipofile a.cpp b.cpp c.cpp
```

See [“Using Profile-Guided Optimization: An Example”](#) later in this chapter for a description of how to use multifile IPO with profile information for further optimization.

Creating a Multifile IPO Executable Using a Project Makefile

Most applications use a make file or something similar to call a linker such as `link.exe`. This is done automatically when you compile and link with `icl`. Therefore, when `-Qipo` must result in a separate linking step, you must use the Intel linker driver `xilink.exe` instead, as follows:

```
prompt> xilink -Qipo <LINK_commandline>
```

`-Qipo` enables multifile IPO

`<LINK_commandline>` is your linker command line

Use the `xilink.exe` syntax when you use a makefile instead of step 2 in the example [“Creating a Multifile IPO Executable”](#). The following example places the multifile IPO executable in `filename.exe` :

```
prompt> xilink -Qipo -out:filename.exe a.obj b.obj c.obj
```



NOTE. For a `-Qipo` command line that results in linkage, the mixed order of object files and linker arguments is not preserved. Instead, the specified linker arguments follow the `-Qipo` optimized object and other objects specified in the linker's command line.

Controlling Inline Expansion of User Functions (`-Obn`, `-Qip_no_inlining`)

The compiler enables you to control the amount of inline function expansion, with `-Obn`. See [Table 5-2](#) for details.

Table 5-2 Summary of `-Obn` and `-Qip_no_inlining` Options

Option	Effect
<code>-Ob0</code>	Disables inline expansion of user-defined functions.
<code>-Ob1</code>	Enables inlining of functions declared with the <code>__inline</code> keyword, and enables inlining according to the C++ language. This option is the default if <code>-O1</code> , <code>-O2</code> , or <code>-Ox</code> are specified.

Table 5-2 Summary of `-Obn` and `-Qip_no_inlining` Options (Continued)

Option	Effect
<code>-Ob2</code>	Enables inlining of any function. However, the compiler decides which functions are inlined. This option enables interprocedural optimizations and has the same effect as specifying the <code>-Qip</code> option.
<code>-Qip_no_inlining</code>	This option is only useful if <code>-Qip</code> or <code>-Ob2</code> is also specified. In this case, <code>-Qip_no_inlining</code> disables inlining that would result from the <code>-Qip</code> or <code>-Ob2</code> interprocedural optimizations, but has no effect on other interprocedural optimizations. This option has no effect on user-directed inlining (<code>-Ob1</code>).

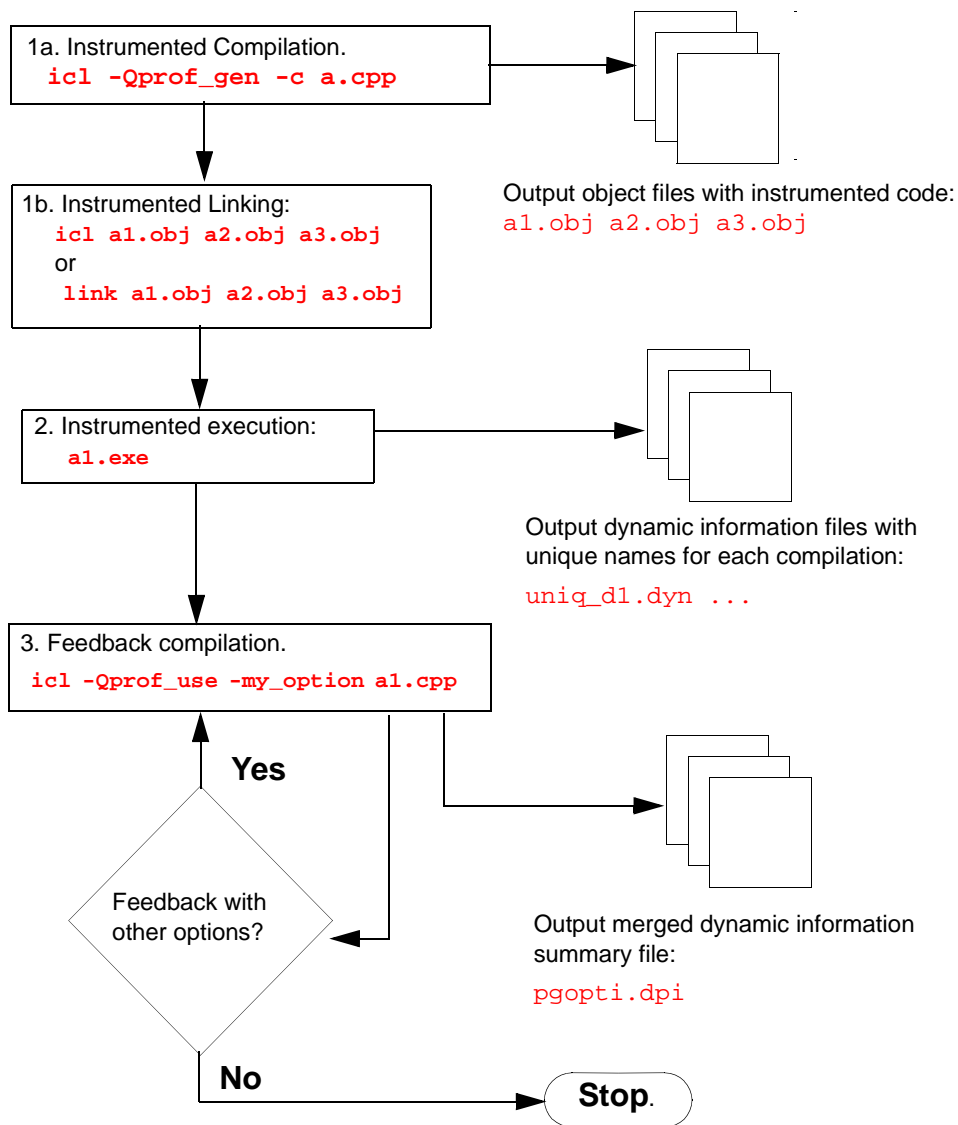
Profile-Guided Optimization (PGO): Three Phases

Profile-guided optimizations tell the compiler which areas of an application are most frequently executed. By knowing these areas, the compiler is able to be more selective and specific in optimizing the application. For example, the use of PGO often allows the compiler to make better decisions about function inlining, thereby increasing the effectiveness of interprocedural optimizations. The PGO methodology requires three phases:

1. *instrumentation compilation* and linking with `-Qprof_gen`
2. *instrumented execution* by running the executable
3. *feedback compilation* with `-Qprof_use`

These steps are summarized in [Figure 5-1](#), followed by option descriptions and an example. A key factor in deciding whether you want to use PGO lies in knowing which sections of your code are the most heavily used. If the data set provided to your program is very consistent and it elicits a similar behavior on every execution, then PGO can probably help optimize your program execution. However, different data sets can elicit different algorithms to be called. This can cause the behavior of your program to vary from one execution to the next.

Figure 5-1 Phases of Basic Profile Guided Optimization



In cases where your code behavior differs greatly between executions, you have to ensure that the benefit of the profile information is worth the effort required to maintain up-to-date profiles. Only two options are used in a basic profile-guided optimization. These options are `-Qprof_gen` and `-Qprof_use`, used in phases 1 and 3 as described in [Table 5-3](#).

Table 5-3 Basic Profile-Guided Optimization Options

Option	Description
<code>-Qprof_gen</code>	Instructs the compiler to produce instrumented code in your object files in preparation for instrumented execution. <i>Note:</i> The dynamic information files are produced in phase 2 when you run the executable.
<code>-Qprof_use</code>	Instructs the compiler produce a profile-optimized executable and merges available dynamic information (<code>.dyn</code>) files into a <code>pgopti.dpi</code> file. If you perform multiple executions of the instrumented program, <code>-Qprof_use</code> merges the dynamic information files again and overwrites the previous <code>pgopti.dpi</code> file.

Basic PGO Options and Environment Variables

[Table 5-4](#) describes environment values to determine the directory in which to store dynamic information files or whether to overwrite `pgopti.dpi`. Refer to your operating system documentation for instructions on how to specify environment values.

Table 5-4 Profile-Guided Optimization Environment Variables

Variable	Description
<code>PROF_DIR</code>	Specifies the directory in which dynamic information files are created. This variable applies to all three phases of the profiling process.
<code>PROF_NO_CLOBBER</code>	Alters the feedback compilation phase slightly. By default, during the feedback compilation phase, the compiler merges the data from all dynamic information files and creates a new <code>pgopti.dpi</code> file, even if one already exists. When this variable is set, the compiler does not overwrite the existing <code>pgopti.dpi</code> file. Instead, the compiler issues a warning and you must remove the <code>pgopti.dpi</code> file if you want to use additional dynamic information files.

Using Profile-Guided Optimization: An Example

The following is an example of the basic PGO phases:

1. **Instrumentation Compilation and Linking**—Use `-Qprof_gen` to produce an executable with instrumented information; for example:

```
prompt> icl -Qprof_gen -c a1.cpp a2.cpp a3.cpp
prompt> icl a1.obj a2.obj a3.obj
```

In place of the second command, you could use the linker (`link.exe`) directly to produce the instrumented program. If you do this, make sure you link with the `libirc.lib` library.

2. **Instrumented Execution**—Run your instrumented program with a representative set of data to create a dynamic information file.

```
prompt> a1.exe
```

The resulting dynamic information file has a unique name and `.dyn` suffix every time you run `a1.exe`. The instrumented file helps predict how the program runs with a particular set of data. You can run the program more than once with different input data.

3. **Feedback Compilation**—Compile and link the source files with `-Qprof_use` to use the dynamic information to optimize your program according to its profile:

```
prompt> icl -Qprof_use -Qip a1.cpp a2.cpp a3.cpp
```

Besides the optimization, the compiler produces a `pgopti.dpi` file. You typically specify the default optimizations (`-O2`) for phase 1, and specify more advanced optimizations (`-Qip` or `-Qipo`) for phase 3. This example used `-O2` in phase 1 and the `-Qip` in phase 3.



NOTE. The compiler ignores the `-Qip` or the `-Qipo` options with `-Qprof_gen`.

Guidelines for Using PGO

When you use PGO, consider the following guidelines:

- Minimize the changes to your program after instrumented execution and before feedback compilation. During feedback compilation, the compiler ignores dynamic information for functions modified after that information was generated.



NOTE. *The compiler issues a warning that the dynamic information corresponds to a modified function.*

- Repeat the instrumentation compilation if you make many changes to your source files after execution and before feedback compilation.

Profile Guided Optimizations Using a Function Order List

A function order list is a text file that specifies the order in which the linker should link the non-static functions of your program. This improves the performance of your program by reducing paging and improving code locality. Profile-guided optimizations support the generation of a function order list to be used by the Microsoft Visual C++ Linker. The compiler determines the order using profile information.

To enable the Intel C/C++ compiler and `proforder` tool to generate a function order list, you must use the `-Qprof_genx` and `-Qprof_dir` described in [Table 5-5](#).

Table 5-5 **Advanced Profile-Guided Optimization Options**

Option	Description
<code>-Qprof_genx</code>	Generates an instrumented object file and creates a static profile information file (<code>.spi</code>). When you use <code>-Qprof_genx</code> instead of <code>-Qprof_gen</code> , you are able to use the <code>proforder</code> tool to create a function order list for the linker. However, using <code>-Qprof_genx</code> for instrumentation requires more memory at runtime, and it produces larger <code>.dyn</code> files and you cannot execute parallel make files when you use <code>-Qprof_genx</code>

Table 5-5 Advanced Profile-Guided Optimization Options (Continued)

Option	Description
<code>-Qprof_dir <dirname></code>	Specifies the directory where <code>.dyn</code> files are to be created. The default is the directory where the program is compiled. The specified directory must already exist. You should specify the same <code>-Qprof_dir</code> option for both the instrumentation and feedback compilations. If you move the <code>.dyn</code> files, you need to specify the new path.

You will need to use the utilities `profmerge` and `proforder` described later in this chapter in [“Utilities for Profile-Guided Optimization”](#).

Function Order List Usage Guidelines

Use the following guidelines to create a function order list.

1. The order list only affects the order of non-static functions.
2. Do not use `-Qprof_genx` to compile two files from the same program simultaneously. This means that you cannot use the `-Qprof_genx` option with parallel `makefile` utilities.
3. You must compile with `-Gy` to enable function-level linking. This option is active when you specify either `-O1` or `-O2`.

Function Order List Example

Assume you have a C program that consists of files `file1.c` and `file2.c` and that you have created a directory for the profile data files in `c:\profdata`. Do the following to generate and use a function order list.

1. Compile your program by specifying `-Qprof_genx` and `-Qprof_dir`:

```
prompt> icl -FeMYPROG -Qprof_genx -Qprof_dir
c:\profdata file1.c file2.c
```
2. Run the instrumented program on one or more sets of input data.

```
prompt> MYPROG.exe
```

The program produces a `.dyn` file each time it is executed.
3. Merge the data from one or more runs of the instrumented program using the `profmerge` tool to produce the `pgopti.dpi` file.

```
prompt> profmerge -Qprof_dir c:\profdata
```

4. Generate the function order list using the `proforder` tool. By default, the function order list is produced in the file `proford.txt`.

```
prompt> proforder -Qprof_dir c:\profdata
-o MYPROG.txt
```

5. Compile your application with profile feedback by specifying the `-Qprof_use` and the `/ORDER` option to the linker. Again, use the `-Qprof_dir` option to specify the location of the profile files.

```
prompt> icl -FeMYPROG -Qprof_use -Qprof_dir
c:\profdata file1.c file2.c -link /ORDER:@MYPROG.txt
```

Utilities for Profile-Guided Optimization

Use the `profmerge` and `proforder` to help you merge dynamic information files and improve the performance of your programs.

The `profmerge` Utility

This tool merges the dynamic profile information files (`.dyn`). The compiler executes this tool automatically during the feedback compilation phase when you specify `-Qprof_use`. You need this tool if you are generating a function order list for use with the `/ORDER` option provided by the Microsoft Visual C++ linker. The command-line usage for `profmerge` is as follows:

```
prompt> profmerge [-nologo] [-prof_dir <dir_name>]
```

This merges all `.dyn` files in the current directory or the directory specified by `-prof_dir`, and produces the summary file `pgopti.dpi`.

The `proforder` Utility

Generates a function order list for use with the `/ORDER` option provided by the linker. The command line syntax for this tool is as follows:

```
prompt> proforder [-nologo] [-prof_dir <dir_name>]
[-o <order_file>]
```

`<dir_name>` is the directory containing the profile files (`.dpi`, `.dyn`, and `.spi`)

`<order_file>` is the optional name of the function order list file. The default name is `proford.txt`.

The `proforder` tool is used as part of the feedback compilation phase, in conjunction with the Microsoft Visual C++ linker, to improve the performance of your program.

Function Call to Dump Profile Data Explicitly

As part of the instrumented execution phase of profile-guided optimization, the instrumented program writes profile data to the dynamic information file (`.dyn file`). The file is written after the instrumented program returns normally from `main()` or calls the standard C `exit` function. For programs that do not terminate normally, the `_PGOPTI_Prof_Dump` function is provided. During the instrumentation compilation (`-Qprof_gen`) you can add a call to this function to your program. You should add the following function prototype prior to the call:

```
void _cdecl _PGOPTI_Prof_Dump(void);
```



NOTE. *You must remove the call prior to the feedback compilation (`-Qprof_use`). You can also enable and disable the call by commenting it out depending on whether you want the dynamic information file to be overwritten during feedback compilation.*

Specifying Compilation Output

6

This chapter describes all the Intel C/C++ options that determine the output created by the compiler. By default, the compiler converts source code directly to an executable file. However, with certain options, you can control the output file that is produced by the compiler.

By having control of the output file, you can create a file at any of the compilation phases such as assembly, object, or executable. If no errors occur during processing, you can use the output files from a particular phase as input to a later compiler invocation. [Table 6-1](#) describes the options to control the output.

Table 6-1 **Compiler Input and Output Summary**

Last Phase Completed	Option	Compiler Input	Compiler Output
Preprocessing	-P, -E, or -EP	source files	preprocessed files (See “Preprocessing Only (-E, -EP, and -P)” in Chapter 7 for more information about these options.)
Syntax checking	-Zs	C or C++ source files preprocessed files	diagnostic list
Compilation	-S	source files preprocessed files	assembly-language files
Assembly	-c	source files preprocessed files assembly-language files	unlinked object files

continued

Table 6-1 Compiler Input and Output Summary (Continued)

Last Phase Completed	Option	Compiler Input	Compiler Output
Linking	(default)	source files preprocessed files assembly files object files libraries	executable file map file linkable object file
compilation, linking, or assembly	-Fa, -Fo, -Fe	source, assembly, or object files	assigns a name to an output file of your choosing

Parsing for Syntax Only (-Zs)

Use the **-Zs** option to stop processing source files after they have been parsed for C/C++ language errors. This option provides a method to quickly check whether sources are syntactically and semantically correct. The compiler creates no output file. In the following example, the compiler checks a file named `progl.cpp`.

Any diagnostics appear on the standard error output.

```
prompt> icl -Zs progl.cpp
```

Producing an Assembly Code Listing (-S)

Use the **-S** option to generate assembly files with the `.asm` suffix. The compiler does not perform the assembly and linking phases after the source files have been preprocessed and compiled. The file contains comments indicating the source-code line and column that corresponds to each assembly language instruction.

Use the **-Fa** option to name the resulting file. By default, the compiler uses the name of each source file with the `.asm` extension.

For example, the following command creates two assembly language files, `file.asm` and `file2.asm`, that can later be assembled and linked:

```
prompt> icl -S file.cpp file2.cpp
```

You can assemble the `.asm` files to produce `file.obj` and `file2.obj`:

```
prompt> icl -c file.asm file2.asm
```

Assembly files produced by the Intel C/C++ Compiler can be assembled using MASM 6.11 or higher. Certain additional options are required when using MASM 6.11 to assemble files produced by the Intel C/C++ Compiler. Use a command line similar to the following to call MASM:

```
ml -c -coff -Cp file.asm
```



CAUTION. Any existing file with the same name and extension is overwritten.

The following is an example of a portion of an assembly file listing:

```
$B1$3:      ; 100      ; preds: B1$2 B1$3
          add    eax, 4                ; 7.5
          jne    $B1$3      ; PROB 95%          ; 8.5
                                   ; LOE: eax, ebx, esi, edi, esp
```

- **\$B1\$3** identifies the beginning of the third basic block in the first function of the file. A basic block is a set of instructions with the property that if the first instruction is executed then all of the subsequent instructions in the set are also executed.
- **; n** following the basic block label is the block execution count. This count is only printed when the `-Qprof_use` option is used. It indicates the average number of times a block was executed when the instrumented program was run. See [“Basic Profile-Guided Optimization Options” in Chapter 5](#) for more information on `-Qprof_use`.
- **; Preds** is a list of predecessors of the current basic block. Predecessors are blocks that can transfer control to the current basic block.
- The numbers (**x.y**) following the semicolon (**;**) at the end of each instruction indicate the source line number and column corresponding to that assembly language instruction.
- **; Prob 95%** indicates the probability assigned to a jump.

- ; LOE indicates a list of registers which are live on exit from the current basic block. These are registers that contain values to be used by succeeding basic blocks.

Suppressing Linking (-c)

Use the `-c` option to suppress linking. For example, entering the following command produces the object files `file.obj` and `file2.obj`:

```
prompt> icl -c file.cpp file2.cpp
```



NOTE. *The preceding command does not link these files to produce an executable file.*

Using the Microsoft Assembler to Produce Object Code (-Quse_asm)

Use the `-Quse_asm` option to generate assembly code from input source files and then call the Microsoft assembler (MASM) to generate the object files. By default, the compiler generates an object file directly without going through the assembler.

For example, the following command generates assembly code, then calls the assembler to generate object files:

```
prompt> icl -c -Quse_asm file1.cpp
```



NOTE. *MASM does not support all the features used on the object files for C++. Therefore you cannot assemble all files generated by `icl` with MASM. See the MASM documentation for details.*

MASM 6.11 handles a maximum of approximately 32760 lines. If you attempt to assemble a `.asm` file with more than 32760 lines, the assembler aborts with a fatal error. Also, if you compile with the `-Quse_asm` option and the intermediate assembly file exceeds 32760 lines, the assembler aborts.

Assembly files are produced by the `-S` option and are implicitly produced and assembled by the `-Quse_asm` option. Due to limitations in MASM and the assembly file format, it may not be possible to assemble `.asm` files produced by the Intel C/C++ Compiler. The following are specific features that are incompatible with assembly file usage:

- Thread Local Storage using `__declspec(thread)` syntax
- Debug Information Generation using the `-Zi`, `-Z7`, or `-Zd` options.
- `#pragma comment` and the embedding of default library directives in the object file. The following options are affected: `-LD`, `-LdD`, `-MD`, `-MDd`, `-ML`, `-MLd`, `-MT`, `-MTd`, `-V`
- `COMDAT` generation using the `-Gf` or `-Gy` options is not supported in assembly files.
- Building of multifile C++ programs is not possible when assembly files are used, due to lack of support for `COMDAT` generation.

Linking

If you do not specify any of the preceding phase limitation options, the compiler defaults to creating executable files; providing that your code is error-free. The following sections describe options that you can use to change the name of the output file and how to prepare for debugging.

Naming the Output File (-Fe, -Fo, -Fa)

When compiling and linking a set of source files, you can use the `-Fe`, `-Fo`, or `-Fa` options to give the resulting file a name other than that of the first source or object file on the command line. If you are processing a single file, you can use the `-Fename`, `-Foname`, and `-Faname` options to specify alternate names respectively for executable, object, and assembly files. Do not enter a space between the option and the `name` argument. You can also use these options to override the default filename extensions `.asm` and `.obj`. In the following example, the `-Fo` option assigns the name `file.obj` to an output object file rather than the default (`x.obj`). The `-c` option directs the compiler to suppress linking.

```
prompt> icl -c -Fofile.obj x.cpp
```

In the following example, the command produces an executable file named `outfile.exe` as the result of compiling and linking two source files.

```
prompt> icl -Feoutfile.exe file.cpp file2.cpp
```

When compiling one or more files, the `-Fe`, `-Fa`, and `-Fo` options can be given a *dirname* (that is, a directory name) argument. The *dirname* argument must end in a forward slash “/” or backslash “\” character, and it must name an existing directory. In this case, the compiler uses the default convention in naming the executable, assembly, or object files produced, but the files will be placed in the directory specified by *dirname*.

In the following example, assume that `obj_dir` is an existing directory. The `-Fe` option causes the compiler to create the executable `myprog.exe` in the current directory. The `-Fo` option causes the compiler to create the object files `a.obj`, `b.obj`, and `c.obj` and place them in the directory `obj_dir`.

```
prompt> icl -Femyprog.exe -Foobj_dir/ a.cpp b.cpp c.cpp
```

Do not enter a space between the option and the *dirname* argument. You can specify different name arguments for each of the `-Fe`, `-Fa`, and `-Fo` options. The compiler does not remove objects that it produces, even when the compilation proceeds to the link phase.

Preparing for Debugging (-Zi, -Oy, -Oy-)

Use the `-Zi` option to direct the compiler to generate code to support symbolic debugging. For example:

```
prompt> icl -Zi prog1.cpp
```

The compiler does not support the generation of debugging information in assembly files. If you specify the `-Zi` option with `-Quse_asm`, the resulting object file will not contain debugging information. If you specify the `-Zi` option with `-S` and later assemble the resulting assembly file, the resulting object file will not contain debugging information. If you specify the `-Zi` option with `-Fa`, the resulting object file will contain debugging information, but the assembly file will not.

Also, because the Intel C/C++ Compiler does not use Microsoft's PDB format for debug information, object files with debug information are larger than those produced by the Microsoft Visual C++ Compiler.

The compiler uses `-Od` as the default when you specify `-Zi`. Specifying the `-Zi` or `-Od` option automatically disables the `-Oy` option.

The `-Oy` option, which is enabled by default or when `-O1` or `-O2` is specified, allows the compiler to use the `ebp` register as a general purpose register in optimizations. However, most debuggers expect `ebp` to be used as a stack frame pointer, and cannot produce a stack backtrace unless this is so. The `-Oy-` option instructs the compiler to generate code that maintains and uses `ebp` as a stack frame pointer, without turning off optimization, so that a debugger can still produce a stack backtrace. Using this option reduces the number of available general purpose registers by one, and can result in slightly less efficient code.

Optimizations and Support for Symbolic Debugging

As described in the preceding section, specifying `-Zi` or `-Od` automatically disables the `-Oy` option. The compiler lets you generate code to support symbolic debugging while the `-O1`, or `-O2` optimization options are specified on the command line along with `-Zi`. However, you can receive these unexpected results:

- If you specify the `-O1`, or `-O2` options with the `-Zi` option, some of the debugging information returned may be inaccurate as a side-effect of optimization.
- If you specify the `-O1`, or `-O2` options, the `-Oy` option will not be disabled. In this case, if you want to maintain the frame pointer while generating debug information, you must explicitly specify the `-Oy-` option on the command line to disable `-Oy`.

The effects of using the `-Zi` option with the optimization options are summarized in Table 6-2.

Table 6-2 **Effects of Using -Zi with Optimization Options**

These options	Imply these results
<code>-Zi</code>	debugging information produced, <code>-Od</code> , <code>-Oy</code> disabled
<code>-Zi -O2</code>	debugging information produced, <code>-O2</code> optimizations enabled
<code>-Zi -O2 -Oy-</code>	debugging information produced, <code>-O2</code> optimizations enabled, <code>-Oy</code> disabled
<code>-Zi -Qip</code>	limited debugging information produced, <code>-Qip</code> option enabled

See the respective sections in [Chapter 4, Optimizations](#) for detailed descriptions of the optimizations options listed in Table 6-2.

Preprocessing

7

This chapter describes the options you can use from the command line to direct the operations of the preprocessor. Preprocessing performs such tasks as macro substitution, conditional compilation, and file inclusion. [Table 7-1](#) provides a summary of the preprocessing options.

Table 7-1 **Options to Control Preprocessing**

Option	Description
-QAname[(tokens)]	Associates a symbol name with the specified sequence of values; equivalent to an <code>#assert</code> preprocessing directive.
-QA-, -u	Causes all predefined macros (other than those beginning with <code>__</code>) and assertions to be inactive.
-C	Preserves comments in preprocessed source output.
-Dname[=value]	Defines the macro <i>name</i> and associates it with the specified <i>value</i> . The default (-Dname) defines a macro with a <i>value</i> of 1.
-E	Directs the preprocessor to expand your source module and write the result to standard output.
-EP	Same as -E but does not include <code>#line</code> directives in the output.
-P	Directs the preprocessor to expand your source module and store the result in a file in the current directory.
-Uname	Suppresses any automatic definition for the specified macro.

Preserving Comments in Preprocessed Source Output (-C)

Use the `-C` option to preserve comments in your preprocessed source output.

Preprocessing Only (-E, -EP, and -P)

Use either the `-E` or the `-P` option to preprocess your source files without compiling them.

When you specify the `-E` option, the compiler's preprocessor expands your source module and writes the result to standard output. The preprocessed source contains `#line` directives, which the compiler uses to determine the source file and line number during its next pass. For example, to preprocess two source files and write them to `stdout`, enter the following command:

```
prompt> icl -E prog1.cpp prog2.cpp
```

When you specify the `-P` option, the preprocessor expands your source module and stores the result in a file in the current directory. There is no way to change the default name. The preprocessor uses the name of each source file with the `.i` extension. For example, the following command creates two files named `prog1.i` and `prog2.i`, which you can use as input to another compilation:

```
prompt> icl -P prog1.cpp prog2.cpp
```

The `-EP` option can be used in combination with `-E` or `-P`. It directs the preprocessor to not include `#line` directives in the output. Specifying `-EP` alone is the same as specifying `-E -EP`.



CAUTION. When you use the `-P` option, any existing files with the same name and extension are overwritten.

Defining Macros (-QA, -QA-, -u, -D, and -U)

You can use the `-QA` and `-D` options to define the assertion and macro names to be used during preprocessing. The `-U` option directs the preprocessor to suppress an automatic definition of a macro.

Use the `-QA` option to make an assertion. This option performs the same function as the `#assert` preprocessor directive. The form of this option is:

`-QAname[(value)]`

name Indicates an identifier for the assertion.

value Indicates a value for the assertion. If a value is specified, it should be quoted, along with the parentheses delimiting it.

For example, to make an assertion for the identifier `fruit` with the value `orange,banana`, use the following command:

```
prompt> icl -QA"fruit(orange,banana)" prog1.cpp
```

The compiler provides a number of predefined macros. For a list of predefined macros available to the Intel C/C++ Compiler, see [“Predefined Macros”](#).

Enter `-QA-` or `-u` to suppress all predefined macros, except for those beginning with the double underscore.

Use the `-D` option to define a macro. This option performs the same function as the `#define` preprocessor directive. The form of this option is:

`-Dname[=value]`

name The name of the macro to define.

value Indicates a value to be substituted for *name*. If you do not enter a value, *name* is set to 1. The value should be quoted if it contains non-alphanumerics.

For example, to define a macro called `SIZE` with the value 100 use the following command:

```
prompt> icl -DSIZE=100 prog1.cpp
```

Use the `-Uname` option to suppress any automatic definition for the specified name. The `-U` option performs the same function as a `#undef` preprocessor directive.

For more details about preprocessor directives, see a language reference such as *C: A Reference Manual*.

Predefined Macros

The predefined macros available for the Intel C/C++ Compiler are described in [Table 7-2](#). The **Default** column describes whether the macro is enabled (ON) or disabled (OFF) by default. The **Disable** column lists the option that disables the macro; no indicates that the macro cannot be disabled.

Table 7-2 Predefined Macros

Macro Name	Default	Disable	Description / When Used
<code>__ICL=n</code>	n=400	no	Identifies the Intel C/C++ Compiler for Win32 systems Version 3.0 (n=300).
<code>__cplusplus</code>	C++ only	no	Defined when compiling C++ source.
<code>_WIN32</code>	ON	-u	Defined for Win32 applications
<code>_MSC_EXTENSIONS</code>	ON	-u	Specifies Microsoft Visual C++ language extensions.
<code>_MSC_VER=n</code>	n=1100	-u	Defined for Microsoft Visual C++ compatibility. Set by the installation program through options in the <code>icl.cfg</code> file. The value of <code>n</code> is based on the version of Visual C++ installed:
			<u>version</u> <u>n=</u>
			4.0 1000
			4.1 1010
			4.2 1020
			5.0 1100
			6.0 1200

continued

Table 7-2 **Predefined Macros (Continued)**

Macro Name	Default	Disable	Description / When Used
<code>_M_IX86=n</code>	ON, $n=600$	<code>-u</code>	defined based on the processor option you specify: $n=300$ if you specify the <code>-GB</code> or <code>-G3</code> option $n=400$ if you specify the <code>-G4</code> option $n=500$ if you specify the <code>-G5</code> option $n=600$ if you specify the <code>-G6</code> option
<code>_DLL</code>	OFF	<code>-u</code>	defined if you specify the <code>-MD</code> option
<code>_MT</code>	OFF	<code>-u</code>	defined if you specify the <code>-MD</code> , <code>-MT</code> , or <code>-LD</code> option
<code>_CHAR_UNSIGNED</code>	OFF	<code>-u</code>	defined if you specify the <code>-J</code> option
<code>_CPPRTTI</code>	OFF	<code>-u</code>	defined if you specify the <code>-GR</code> option for C++ only
<code>_CPPUNWIND</code>	OFF	<code>-u</code>	defined if you specify the <code>-GX</code> option for C++ only

See “[Predefined Macros for Standard Conformance](#)” in [Chapter 8](#) for the list of predefined macros required for ANSI conformance standard.

Printing Include-file Dependencies (-QH)

Use the `-QH` option to print the pathname for each file compiled into the source with a `#include` directive. Pathnames can be absolute or relative. The compiler displays the dependencies on the standard output and prints the path for each included file on a separate line. The compilation process stops after preprocessing is completed.

In the following example, the source file, `dtest.cpp`, includes three other files. Enter the following command to display the dependent files:

```
prompt> icl -QH dtest.cpp
./d1.h
./d2.h
./d3.h
```

Printing Makefile Dependencies (-QM)

Use the `-QM` option to generate list of makefile dependency lines for each source file in the compilation. The compiler displays the dependency lines based on `#include` lines that appear in each source file. For example, the output of a simple program with one include file might be as follows:

```
hello.obj: hello.c  
hello.obj: c:/Msdev/Include/stdio.h
```

C/C++ *Language Features*

8

This chapter describes the C and C++ language implementation of the Intel C/C++ Compiler. This chapter does not teach you how to program in C or C++. Instead, it covers the following topics:

- conformance to the ANSI/ISO standard for C
- options that support the C++ language

Conformance to C Standards

The compiler accepts two C language dialects: strict ANSI conformance, and extended. You can select the style that best suits your application.

The compiler's implementation of C conforms to the ANSI/ISO standard (ISO/IEC 9899:1990) for the C language. The standard specifies that a conforming implementation of a C compiler must meet minimum requirements for certain translation limits. In all cases, the compiler exceeds these limits. [Appendix A, "Compiler Limits,"](#) lists the tested values. The compiler also accepts extensions to the standard.

By default, the compiler does not flag all extensions with error or warning messages. You must use the `-Za` option for the compiler to strictly enforce the standard. For more information on this option, see ["Strict ANSI Dialect \(-Za\)"](#) in this chapter.

C Language Dialects

The compiler supports two C language dialects: strict ANSI, and extended. By default, the compiler accepts the extended dialect. You can use the `-Za` option to select the language as defined by the ANSI/ISO standard (ISO/IEC 9899:1990) with no deviation.

The following sections describe each dialect. See *C: A Reference Manual*, or *The C Programming Language*, all listed in the “About This Manual” section, for a full description of the C language.

Strict ANSI Dialect (-Za)

Use the `-Za` option to select the strict ANSI dialect. The compiler does not deviate from the standard when this option is in effect. While following strict conformance to the standard, the compiler flags any nonstandard code in the source file with warnings or error messages.

Extended Dialect (-Ze)

The compiler accepts the language specified in the standard when you enter the `-Ze` option on the invocation line, with extensions in the areas described below.

Extensions for files and data storage:

- The input file can contain no text, so you can use an empty file as input to the compiler.
- The last member of a structure can have an incomplete array type. However, that member cannot be the only member of the structure, otherwise, the structure size would be zero.
- A file-scope array can have an incomplete `struct` or `union` type as its element type. The `struct` or `union` type must be completed before the array is subscripted and by the end of the compilation if the array is defined in the compilation.
- You can declare an `enum` tag name, and then define it later in the source file.
- An initializer expression that is a single value, and is used to initialize an entire static array, structure, or union, need not be enclosed in braces. Standard C requires the braces.

- The compiler generates a remark message for a storage-class specifier appearing anywhere other than the first position in a list of specifiers, as in `int static`.

Extensions for pointers:

- In an initializer, you can cast a pointer constant value to an integral type if the integral type is big enough to contain it.
- In an integral constant expression, you can cast an integer constant to a pointer type and then back to an integral type.
- The compiler accepts assignments of pointers to integers and to other incompatible pointer types without an explicit cast.
- You can select a field in the form `p->field`, even if `p` does not point to a `struct` or `union` that contains `field`. The `p` variable must be a pointer. Likewise, `x.field` is allowed even if `x` is not a structure or union that contains `field`. The `x` variable must be an `lvalue`. For both cases, all definitions of `field` as a field must have the same type and offset within their structure or union.

Extensions for types and syntax:

- Bit fields can have `enum` base types, or integral types other than `int` and `unsigned int`.
- You can use `long float` as a synonym for `double`.
- You can place arbitrary text at the end of preprocessing directives.
- Numbers are scanned according to the syntax for numbers rather than the `pp-number` syntax. Thus, `0x123e+1` is scanned as three tokens instead of one invalid token. Under strict ANSI conformance, the compiler uses the `pp-number` syntax. See the standard text for more information on the `pp-number` syntax.

Extensions for predicates:

- You can use the `#assert` and `#unassert` macros to define and test predicate names.

Extensions with Warnings. When you use the `-Ze` option, the compiler accepts and tags the extensions in syntax and semantics as described below; however, these extensions are tagged with a warning message.

Extensions for syntax with warnings:

- You can have an extra comma at the end of an `enum` list.
- You can omit the final semicolon preceding the closing brace `}` of a structure or union.
- A right brace can immediately follow a label definition. Normally, a statement must follow a label definition.
- You can have an empty declaration (a semicolon with nothing before it).

Extensions for semantics with warnings:

- You are allowed assignment and pointer differences between pointers to types that are interchangeable but not identical, such as `unsigned char *` and `char *`. However, the compiler issues a warning. Assignment of a string constant to a pointer to any kind of character is allowed without a warning.
- You can compare a pointer to a `void` to a pointer of another kind using the `>`, `>=`, `<`, or `<=` operators without using an explicit type cast. Standard C allows such comparisons using `==` or `!=` and issues no warnings.
- You can insert inline assembly code using the `asm` keyword. In the strict ANSI C dialect, such insertions are disabled.
- You can have freestanding tag declarations in the parameter declaration list for a function with old-style parameters.
- The compiler generates a warning if an overflow is detected while folding signed integer operations on constants.

Predefined Macros for Standard Conformance

The standard requires that certain predefined macros be supplied with the compiler. [Table 8-1](#) lists the macros the compiler defines in accordance with the standard.

Table 8-1 **Predefined Macros for Standard Conformance**

Macro	Description
<code>__DATE__</code>	The date of compilation as a string literal in the form "Mmm dd yyyy".
<code>__FILE__</code>	A string literal representing the name of the file being compiled.
<code>__LINE__</code>	The current line number as a decimal constant.
<code>__STDC__</code>	The constant 1 under the ANSI conformance dialect (-Za); not defined in the extended dialect.
<code>__TIME__</code>	The time of compilation, as a string literal in the form "hh:mm:ss".

The compiler provides some predefined macros in addition to the predefined macros required by the standard. For a list of these macros see [“Predefined Macros” in Chapter 7](#) of this manual and the section on the preprocessor category of the Microsoft *Visual C++ User’s Guide*.

Conformance to C++ Standards

The Intel C/C++ Compiler conforms to the Microsoft implementation of the C++ language. For more information on the Microsoft implementation of the C++ language, see the Microsoft Visual C++ documentation.

Microsoft Compatibility

9

The Intel C/C++ Compiler supports the Microsoft Visual C++ extensions to the C and C++ languages. This chapter describes the differences in the way the Intel C/C++ Compiler interprets some of the Microsoft Visual C++ extensions. The differences are in the following areas:

- Compiler Pragmas
- Microsoft compatibility option (`-Qms`)
- Microsoft Version Compatibility Options (`-Qvcn`)
- Unsupported Compiler Options
- Differences in PCH Support
- Compilation and Execution Differences
- Evaluation of Left Shift Operations
- Use of Friend Injection
- Function Declaration in Scope of Function Defined in a Namespace
- Enum Bit-Field Signedness

Compiler Pragmas

The Intel C/C++ Compiler supports the Microsoft Visual C++ pragmas with the following limitations:

- The pragma `pragma optimize` accepts a list that allows the enabling and disabling of specific optimizations. The Intel compiler ignores this list. The pragma either enables or disables all optimizations specified by options on the command line.

- The following pragmas are accepted without error but have no effect:

```
pragma auto_inline      pragma component
pragma function         pragma include_alias
pragma inline_depth     pragma inline_recursion
pragma intrinsic        pragma setlocale
pragma warning
```

For more information on pragmas, see the on-line help or the documentation that accompanies your copy of Microsoft Visual C++.

Microsoft Compatibility Option (-Qms)

In a limited number of cases, the Microsoft compiler compiles source code without generating an error while the Intel compiler generates an error for the same source code. If this occurs in a file that is part of a third party distribution (such as include files or C++ class library files), it might not be convenient to correct the error in the source file. In such cases, using the `-Qms` option with the Intel compiler might allow a successful compilation.

Microsoft Version Compatibility Options (-Qvcn)

You can use the Intel C/C++ Compiler with Microsoft Visual C++ Versions 4.2 or higher. Some language features are available only in Visual C++ Version 5.0 or later and other features are available only in Version 6.0. The `-Qvcn` option tells the compiler what version of Visual C++ you are using. [Table 9-1](#) lists the valid `-Qvcn` options for each version:

Table 9-1 Microsoft Version Compatibility Options

Option	Compatibility with
<code>-Qvc4</code>	Visual C++ Version 4.x
<code>-Qvc5</code>	Visual C++ Version 5.0
<code>-Qvc6</code>	Visual C++ Version 6.0



NOTE. *The Intel C/C++ Compiler installation program automatically adds the appropriate `-Qvcn` option for your version of Visual C/C++ (if needed) to the Intel C/C++ Compiler configuration file (`icl.cfg`). Therefore, you will typically not need to use the `-Qvcn` option.*

Unsupported Compiler Options

The Intel C/C++ Compiler supports most of the same options as the Microsoft Visual C++ Compiler, as well as Intel-specific options. However, a small subset of Microsoft Visual C++ Compiler options are not supported by the Intel C/C++ Compiler. Most of the unsupported options, while useful for development purposes, are not required to build a working application. The unsupported options are listed in [Table 9-2](#).

Table 9-2 List of Unsupported Microsoft Visual C++ Compiler Options

Option	Description
<code>-Fd</code>	name the PDB file used for debug information for specified source files
<code>-GD</code>	optimize for Windows DLL
<code>-Gi</code>	enable incremental compilation
<code>-GI</code>	edit and continue debugging (same effect as the <code>-Zi</code> option)
<code>-Gm</code>	enable minimal rebuild
<code>-Oa</code>	assume no aliasing
<code>-Ow</code>	assume no cross-function aliasing
<code>-WX</code>	treat warnings as errors
<code>-Yd</code>	put debug information in every object (Microsoft PCH-specific option)
<code>-Zg</code>	generate function prototypes
<code>-Zm</code>	set compiler's memory allocation limit

The Intel C/C++ Compiler issues a remark stating lack of support for many of these options, but it silently ignores the following options: `-Fd`, `-Gi`, `-Gm`, `-Yd`, and `-Zm`.

Differences in PCH Support

There are some differences in how precompiled header (PCH) files are supported between the Intel C/C++ Compiler and the Microsoft Visual C++ Compiler. These differences include the following:

- The PCH information generated by the Intel C/C++ Compiler is not compatible with the PCH information generated by the Microsoft Visual C++ Compiler.
- The Intel C/C++ Compiler does not support PCH generation and use in the same translation unit.
- The Intel C/C++ Compiler does not generate PCH information beyond a point where a declaration is seen in the primary translation unit. The Microsoft Visual C++ compiler generates PCH information beyond this point in some situations.

Compilation and Execution Differences

While the Intel C/C++ Compiler is compatible with the Microsoft Visual C++ Compiler, some differences can prevent successful compilation. Also there can be some incompatible generated-code behavior of some source files with the Intel C/C++ Compiler. In most cases, a modification of the user source file enables successful compilation with both the Intel C/C++ Compiler and the Microsoft Visual C++ Compiler. The differences between the compilers are listed as follows:

Inline Assembly Target Labels

When compiled by the Intel C/C++ Compiler, inline assembly target labels of `goto` statements are case sensitive. The Microsoft Visual C++ Compiler treats these labels in a case insensitive manner. For example, the Intel C/C++ Compiler issues an error when compiling the following code:

```
int func(int x)
{
    goto LAB2;
    /*    error: label "LAB2" was referenced
     *    but not defined
     */
    __asm lab2: mov x, 1
    return x;
}
```

However, the Microsoft Visual C++ Compiler accepts the preceding code. As a work-around for the Intel C/C++ Compiler, when a `goto` statement refers to a label defined in inline assembly, you must match the label reference with the label definition in both name and case.

Preprocessor Macro Expansion

The Intel C/C++ Compiler differs from the Microsoft Visual C++ Compiler in the way it expands preprocessor macros that are used within `#include` directives. In some case the code passes a macro as a parameter to another macro that uses the token-concatenation operator. In such a case, the macro that is used as a parameter is not expanded before concatenation. This is demonstrated in the following example:

```
#define D    var
#define Decl(d)    int my ## d ## ;
Decl(D)
```

Both compilers would preprocess the preceding source and produce the following code:

```
int myD;
```

When a similar macro is used in a `#include` directive, the Intel C/C++ Compiler behaves similarly to the preceding example. However, the Microsoft Visual C++ Compiler performs an extra preprocessing scan through the `#include` directive to produce different results. In the following example, the `D` and `F` macros, when used in the `#include` directive, are not expanded by the Intel C/C++ Compiler .

```

#define D    sys
#define F    stat.h

#define INC(d,f)    < ## d ## / ## f ## >

#include INC(D,F)
// Visual C++ Compiler interpretation:
// # include <sys/stat.h>
// Intel C/C++ Compiler interpretation:
// #include <D/F>

```

The Intel C/C++ Compiler issues an error for this code when it fails to find the file called **D/F**. The Microsoft Visual C++ Compiler expands the **D** and **F** macros and accepts this code. The following alternative code achieves the same result but it works with both the Intel C/C++ Compiler and the Microsoft Visual C++ Compiler:

```

#define D    sys
#define F    stat.h

#define CAT5(a,b,c,d,e)    a ## b ## c ## d ## e
#define INC(d,f)            CAT5(<,d,/,f,>)

#include INC(D,F)

```

Evaluation of Left Shift Operations

The Intel C/C++ Compiler differs from the Microsoft Visual C++ Compiler in the evaluation of left shift operations where the right operand, or shift count, is equal to or greater than the size of the left operand expressed in bits. The ANSI C Standard states that the behavior of such left-shift operations is undefined, meaning a program should not expect a certain behavior from these operations. This difference is only evident when both operands of the shift operation are constants. The following example illustrates this difference between the Intel C/C++ Compiler and the Microsoft Visual C++ Compiler:


```
int x;
int y = 1;

void func()
{
    x = 1 << 32;
    // Visual C++ Compiler generates code to set x=0
    // Intel C/C++ Compiler generates code to set x=1
    y = y << 32;
    // Visual C++ Compiler generates code to set x=1
    // Intel C/C++ Compiler generates code to set x=1
}
```

Use of Friend Injection: Not Recommended

Sometimes a class is a member of a `namespace`, and contains a friend declaration of a function not already declared. In such cases, the Intel C/C++ Compiler injects the function declaration into the `namespace` containing the class, even if the class definition is not lexically within the `namespace` definition. Microsoft Visual C++ Version 4.2 apparently "injects" the function declaration into the scope lexically containing the class definition. As a result of this difference, friendship may be bestowed on a different function, and access errors may result.

Since friend "injection" has recently been eliminated as a language feature by the C++ standardization committee, the interpretation of code which relies on the specific behavior of Microsoft Visual C++ Version 4.2. In contrast, Microsoft Visual C++ Version 5.0 does not accept source like that described at the beginning of this paragraph and reports an error diagnostic message semantically equivalent to the Intel C/C++ Compiler. To avoid such behavior, you should avoid the use of friend injection.

Declaration in Scope of Function Defined in a Namespace

In accordance with the C++ language specification, if a function declaration is encountered within a function definition, the function referenced is taken to be another member of the `namespace` of the containing function; regardless of whether the containing function definition is lexically within a `namespace` definition. The Microsoft Visual C++ Compiler takes the referenced function to be a global function (not in any `namespace`). Functions declared in global or `namespace` scopes are interpreted the same way by both the Intel C/C++ Compiler and the Microsoft Visual C++ Compiler.

Enum Bit-Field Signedness

The Intel C/C++ Compiler and the Microsoft Visual C++ Compiler differ in how they attribute signedness to bit fields declared with an `enum` type. Apparently, Microsoft Visual C++ always considers `enum` bit fields to be signed, even if not all values of the `enum` type can be represented by the bit field. On the other hand, the Intel C/C++ Compiler considers an `enum` bit field to be unsigned, unless the `enum` type has at least one `enum` constant with a negative value. In any case, the Intel C/C++ Compiler produces a warning if the bit field is declared with too few bits to represent all the values of the `enum` type.


Debugging Functions with Aligned Stack Frames Using MSVC++ 4.2

When debugging a function with an aligned frame compiled with the Intel C/C++ Compiler with the MSVC++ 4.2 tools, the Visual C++ debugger displays invalid values for the parameters. This is due to the fact that the MSVC++ 4.2 linker does not support the required debug records to represent parameters in aligned frames.

Ordinarily, when entering a function, the `EBP` register points to the bottom of the parameter block. The debugger accesses function parameters using the `EBP` register plus an offset. However, in an aligned frame the `EBP`

register is decremented an arbitrary number of bytes, in order to align the frame, and a different register is used as a pointer to the parameters. See [Figure 9-1](#).

Figure 9-1 Simplified View of Arbitrary Offset in the Stack



The Microsoft debug format supported by Visual C++ Version 4.2 does not allow flexibility in the base register used for parameter variable access, so the Visual C++ Version 4.2 debugger always uses **EBP** to access the parameters. This causes incorrect values to be displayed for the parameters, since they are not located at the expected address.

To view parameters in aligned functions, you can use assembly level debugging to look at the actual location of the parameter. The Intel C/C++ Compiler Version 4.0 uses the **EBX** register for parameter access in a function with an aligned frame. In such a function, parameters are located at the address specified by **EBX** plus an offset. Because of the padding that occurs within the parameter block to align **__m128** parameters, it might be easier to determine the correct offset by examining a reference to the parameter of interest in the disassembled code.

This problem does not affect use of the Intel C/C++ Compiler with the Visual C++ Version 5.0 tools. To use the Intel C/C++ Compiler with Visual C++ Version 5.0, you must install Visual C++ 5.0 and then reinstall the Intel C/C++ Compiler. For more information on the aligned stack frame feature, see the description of the **-Qsfalign** in [“Stack Frame Alignment Options” in Chapter 13](#).

Diagnostic Information

10

This chapter describes the various messages that the compiler produces. These messages include the sign-on message and diagnostic messages for remarks, warnings, or errors. The compiler always displays any diagnostic message, along with the erroneous source line, on the standard output.

This chapter also describes how to control the severity of diagnostic messages.

Disabling the Sign-on Message (-nologo)

The compiler displays the following message each time that you invoke it:

```
Intel C/C++ Compiler Version x.y.z ID
Copyright (C) years Intel Corporation. All rights reserved.
ID                               The unique identification number for this compiler.
x.y.z                           Identifies the version of the compiler.
years                           The years for which the software is copyrighted.
```

If you want to suppress the message, specify the `-nologo` option.

Printing the List of icl Options (-?, -help)

You can print a list of the most useful `icl` options by specifying the `-help` (or `-?`) option to the compiler. To print this list, use this command:

```
prompt> icl -help
```

or

```
prompt> icl -?
```

Diagnostic Messages

Diagnostic messages provide syntactic and semantic information about your source text. Syntactic information can include syntax errors and use of nonstandard C or C++. Semantic information includes such things as unreachable code.

Diagnostic messages can be any of the following: command-line diagnostics, remark messages, warning messages, error messages, or catastrophic error messages.

Command-line diagnostics. These messages report improper or unrecognized command-line options or arguments. Command-line error messages appear on the standard output in the form:

```
icl: Command line error: message  
message           Describes the error.
```

Command-line warning messages appear as follows:

```
icl: Command line warning: message
```

Language diagnostics. These messages describe diagnostics that are reported during the processing of the source file. These diagnostics have the following format:

```
filename(linenum): type [#]nn: message
```

<i>filename</i>	Indicates the name of the source file currently being processed.
<i>linenum</i>	Indicates the source line where the compiler detects the condition.
<i>type</i>	Indicates the severity of the diagnostic message: <i>warning</i> , <i>remark</i> , <i>error</i> , or <i>catastrophic error</i> .
[<i>#</i>] <i>nn</i>	The number assigned to the <i>error</i> (or <i>warning</i>) message. Hard errors or catastrophes are not assigned a number.
<i>message</i>	Describes the diagnostic.

The following is an example of a warning message:

```
tantst.cpp(3): warning #328: Local variable "increment"
never used.
```

The compiler can also display internal error messages on the standard error. If your compilation produces any internal errors, contact your Intel representative. Internal error messages are in the following form:

```
FATAL COMPILER ERROR: message
```

Remark messages. These messages report common but sometimes unconventional use of C or C++. The compiler does not print or display remarks unless you specify level 4 for the `-W` option, as described in [“Suppressing Warning Messages or Enabling Remarks \(-w, -Wn\)”](#) later in this chapter. Remarks do not stop translation or linking. Remarks do not interfere with any output files. The following are some representative remark messages:

```
function declared implicitly
type qualifiers are meaningless in this declaration
controlling expression is constant
```

Warning messages. These messages report legal but questionable use of the language in the program being compiled. The compiler displays warnings by default. You can suppress warning messages by setting the diagnostic level to 0. Warnings do not stop translation or linking. Warnings do not interfere with any output files. The following are some representative warning messages:

```
declaration does not declare anything
pointless comparison of unsigned integer with zero
possible use of = where == was intended
```

Error messages. These messages report syntactic or semantic misuse of C or C++. The compiler always displays error messages. Errors suppress object code for the module containing the error and prevent linking, but they allow parsing to continue to scan for any other errors. Some representative error messages are:

```
missing closing quote
expression must have arithmetic type
expected a ";"
```

Catastrophic errors. These messages indicate environmental problems. Catastrophic error conditions stop translation, assembly, and linking. If a catastrophic error ends compilation, the compiler displays a termination message on standard output. The following are some representative catastrophic error messages:

```
out of memory
could not open temporary file filename
could not open source file filename
```

Suppressing Warning Messages with lint Comments

The UNIX `lint` program attempts to detect features of a C or C++ program that are likely to be bugs, non-portable, or wasteful. The compiler recognizes three `lint`-specific comments: `/*ARGSUSED*/`, `/*NOTREACHED*/`, and `/*VARARGS*/`. Like the `lint` program, the compiler suppresses warnings about certain conditions when you place these comments at specific points in the source.

Suppressing Warning Messages or Enabling Remarks (-w, -Wn)

Use the `-w` or `-Wn` option to suppress warning messages or to enable remarks during the preprocessing and compilation phases. You can enter the option with one of the following arguments:

<code>-W0, -w</code>	Displays error messages only.
<code>-W1, -W2, -W3</code>	Displays warnings and error messages. The compiler uses this level as the default.
<code>-W4</code>	Displays remarks, warnings, and error messages.

For some compilations, you might not want warnings for known and benign characteristics, such as the K&R C constructs in your code. For example, the following command compiles `newprog.cpp` and displays compiler errors, but not warnings:

```
prompt> icl -W0 newprog.cpp
```

Controlling the Severity of Diagnostics (-Qwd, -Qwr, -Qww, -Qwe)

You can control the severity of some of the diagnostics returned by the compiler. The compiler returns two types of diagnostics:

hard errors	Diagnostics issued for code that is definitely wrong or questionable. The severity of a hard error is not configurable. For hard errors, the message number is never printed. Remarks and warnings are never considered hard errors.
soft diagnostics	All other diagnostics (including remarks and warnings). For soft diagnostics, the message number is always printed. The severity of a soft diagnostic is configurable by the options described below.

In the descriptions below, *tag* represents the number associated with the diagnostic. Multiple tags are permitted, separated by commas.

-Qwd[*tag*,...] Disable the soft diagnostics that corresponds to *tag*.

-Qwr[*tag*,...] Override the severity of the soft diagnostics corresponding to *tag* and make it a remark.

-Qww[*tag*,...] Override the severity of the soft diagnostics corresponding to *tag* and make it a warning.

-Qwe[*tag*,...] Override the severity of the soft diagnostics corresponding to *tag* and make it an error.

For example, the following command line disables soft diagnostic 68 during compilation of the file `a.cpp`:

```
prompt> icl -Qwd68 -c a.cpp
```

The following command line changes the severity of soft diagnostics 68 and 152 to remarks during compilation of the file `a.cpp`.

```
prompt> icl -Qwr68,152 -c a.cpp
```

Example. Assume that you have a file `x.cpp` that contains the following lines:

```
/*  
/* This is a comment  
*/  
  
extern i;
```


If you compile this code with warnings enabled (the default), you will receive the following response from the compiler:

```
x.cpp(2): warning #9: nested comment is not allowed
  /* This is a comment
  ^
x.cpp(5): warning #260: explicit type is missing ("int"
assumed)
  extern i;
  ^
```

If you compile the code with the option `-Qwd9`, (to disable warning number 9), you will receive the following response from the compiler:

```
x.cpp(5): warning #260: explicit type is missing ("int"
assumed)
  extern i;
  ^
```

Limiting the Number of Errors Reported (-Qwnum)

Use the `-Qwnum` option to limit the number of error messages displayed before the compiler aborts. By default, if more than 100 errors are displayed, compilation aborts.

`-Qwnum` Limit the number of error diagnostics that will be displayed prior to aborting compilation to *num*. Remarks and warnings do not count towards this limit.

For example, the following command line specifies that if more than 50 error messages are displayed during the compilation of `a.cpp`, compilation aborts.

```
prompt> icl -Qwn50 -c a.cpp
```

Additional Information about the Compilation

The compilation system issues a variety of self-explanatory information messages about the compilation system components and the process that are not described in this manual.

The compiler allows you to use all the standard run-time libraries that are part of Microsoft Visual C++ Version 4.0 or later. You can determine the libraries used by your applications by controlling the directories which the Microsoft linker searches by using the options described in this chapter.

Managing Libraries

The `LIB` environment variable contains a semicolon-separated list of directories in which the Microsoft linker will search for library (`.lib`) files. If you want the linker to search additional libraries, you can add their names to the command line, to a response file, or to the `icl.cfg` file. In each case, the names of these libraries are passed to the linker before the names of the Intel libraries (`libm.lib` or `libm_chk.lib`) and the Microsoft-provided default libraries that the driver always specifies (`oldnames.lib` and `libc.lib`). For more information on adding library names to the response file and the configuration file (`icl.cfg`), see [“Response Files”](#) and [“Configuration Files” in Chapter 3](#).

To specify a library name on the command line, you must first add the library’s path to the `LIB` environment variable. Then, to compile `file.cpp` and link it with the library `mylib.lib` enter the following command:

```
prompt> icl file.cpp mylib.lib
```

The compiler driver, `icl.exe`, passes file names to the Microsoft linker in the following order:

1. the object file
2. any objects or libraries specified on the command line, in a response file, or in a configuration file
3. the `libm.lib` library by default (or, if you specified the `-QIfdiv` option, the `libm_chk.lib` libraries)
4. the Microsoft-provided libraries `libc.lib` and `oldnames.lib`

Default Libraries

The compiler allows you to use all the standard run-time libraries that are part of Microsoft Visual C++ Version 4.x or higher. By default, the compiler automatically expands a number of standard C, C++, and math library functions. For more information, see [“Inline Expansion of Library Functions \(-Oi, -Oi-\)” in Chapter 4](#).

Library Files

The compiler automatically tells the linker to use the following support libraries.

`libm_chk.lib` The `libm.lib` file contains the math library. The
or `libm.lib` `libm_chk.lib` file contains the math library with
support routines for a floating-point division software
patch for certain steppings of the Pentium processor. For
information on these libraries, see [“Enabling the
Floating-Point Division Check \(-QIfdiv\)”](#).

These libraries are supplied with Microsoft Visual C++ Version 4.0 or later.

`libc.lib` The standard C run-time library.

`oldnames.lib` Contains aliases for non-ANSI functions that normally
begin with an underscore in Microsoft libraries.

If you want to link your program with alternate or additional libraries, specify them at the end of `icl` command line. For example, to compile and link `hello.cpp` with `mylib.lib`, use the following command:

```
prompt> icl -Fehello.exe hello.cpp mylib.lib
```

The `mylib.lib` library appears prior to the `libc.lib` library in the command line for the `LINK` linker.

Math Libraries

In the compiler package, you received the Intel math library, `libm.lib`, which contains optimized versions of the math functions in the standard C run-time library. The functions in the library are optimized for program execution speed on the Pentium processor.

To use the optimized math library, the set up process places `libm.lib` in one of the directories specified in the library search path defined by the `LIB` variable. Intel recommends that you keep `libm.lib` in the first directory specified in the path.

Enabling the Floating-Point Division Check (-Qfdiv)

The `-Qfdiv` option enables a software patch for the floating-point division flaw that exists on some steppings of the Pentium processor. This patch ensures that the precision of your floating-point division calculations are correct.



NOTE. This option (`-Qfdiv`) is no longer enabled by default when you specify either the `-G3` or the `-G5` options.

When the `-Qfdiv` option is enabled, the compiler uses `libm_chk.lib` instead of `libm.lib` to link your programs. The `libm_chk.lib` library contains the support routines for the floating-point division software patch for affected math library functions.

The `-QIfdiv-` option disables the software patch for the floating-point division flaw regardless of whatever other options are specified. When you specify `-QIfdiv-`, the compiler uses simple hardware instructions for floating-point division and affected intrinsics. If you specify the `-QIfdiv-` option, the compiler links with `libm.lib`. Similarly, if you choose not to use the special version of the optimized math library, you must specify `-QIfdiv-`. This option is the default if you specify either the `-G4` or the `-G6` options.

Avoiding Incorrect Decoding of Certain Instructions (-QIOf)

Some instructions have 2-byte opcodes, the first byte of which contains 0f. In rare cases, the Pentium processor incorrectly decodes these instructions. Specify the `-QIOf` option to avoid the incorrect decoding of these instructions. The work-around implemented in the Intel C/C++ Compiler avoids generating the susceptible instructions.

Controlling Compiler-Generated Code

12

This chapter describes options and features that let you control the outcome of Intel compiler-generated code without interfering with the way your program runs.

- [“Specifying Structure Tag Alignments \(-Zp\)”](#) describes how you can set the alignment of structures and unions from the command line instead of adding a pragma to your source file.
- [“Allocation of Zero-initialized Variables \(-Qnobss_init\)”](#) describes how you can place variables initialized to zero in the **DATA** section instead of allowing the compiler to place them in the **BSS** section.

Specifying Structure Tag Alignments (-Zp)

You can specify an alignment constraint for structures and unions in two ways: place a pack pragma in your source file or enter the alignment option on the command line. Both specifications change structure tag alignment constraints.

Use the **-Zp** option to determine the alignment constraint for structure declarations. Generally, smaller constraints result in smaller data sections while larger constraints support faster execution.

The form of the `-Zp` option is:

`-Zpnumber`

`number`

The alignment constraint indicated by one of the following values:

1	1 byte.
2	2 bytes.
4	4 bytes.
8	8 bytes.
16	16 bytes.

For example, to specify 2 bytes as the alignment constraint for all the structures and unions in the file `progl.e`, use the following command:

```
prompt> icl -Zp2 progl.e
```

Allocation of Zero-initialized Variables (-Qnobss_init)

Use the `-Qnobss_init` option to place any variables that are explicitly initialized with zeros in the `__DATA` section. By default, variables explicitly initialized with zeros are placed in the `__BSS` section, but some programs require them to be in the `__DATA` section.

Support for MMX™ Technology and the Streaming SIMD Extensions

13

The Pentium III Processor and other processors such as the Pentium with MMX technology and Pentium II processors have characteristics that enable the development of advanced multimedia applications. The Streaming SIMD Extensions and MMX technology intrinsics are coding extensions to make use of the processors' multimedia features. This chapter describes how to code with and effectively use these powerful and time saving instructions.

MMX Technology Intrinsics

The benefit of coding with MMX technology intrinsics is that you can use the syntax of C function calls and C variables instead of hardware registers. This frees you from the managing of registers and assembly programming. Further, the compiler optimizes the instruction scheduling so that your executable runs faster.

The intrinsics are based on a new `__m64` data type to represent the specific contents of an `mmx` register.

The `__m64` data type, however, is not a basic ANSI C data type, and therefore you must observe the following usage restrictions:

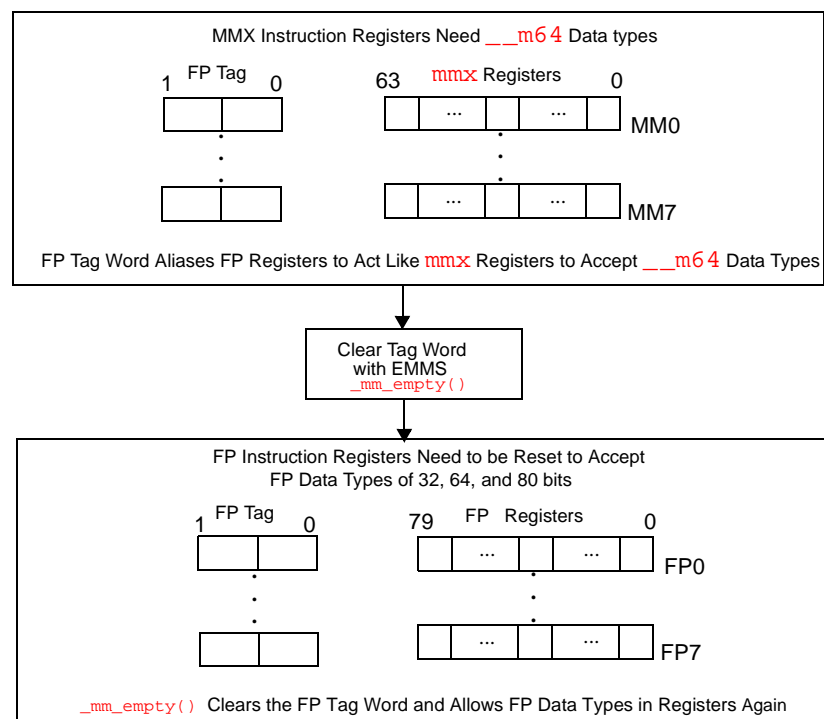
- Use `__m64` data only on the left-hand side of an assignment, as a return value, or as a parameter. You cannot use it with other arithmetic expressions ("`+`", "`>>`", and so on).
- Use `__m64` objects in aggregates, such as unions to access the byte elements and structures; the address of an `__m64` object may be taken.
- Use `__m64` data only with the MMX intrinsics described in this guide.

For complete details of the hardware instructions, see the *Intel Architecture MMX Technology Programmer's Reference Manual*. For descriptions of data types, see the *Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference Manual*.

The EMMS Instruction: Why You Need it and When to Use it

Using **EMMS** is like emptying a container to accommodate new content. For instance, MMX instructions automatically enable an FP tag word in the register to enable use of the `__m64` datatype. This resets the FP register set to alias it as the `mmx` register set. To enable the FP register set again, reset the register state with the **EMMS** instruction or via the `_mm_empty()` intrinsic.

Figure 13-1 Why You Need EMMS to Reset After an MMX™ Instruction





CAUTION. Failure to reset the tag word for FP instructions after using an MMX instruction can result in unexpected execution or poor performance.

Guidelines for When to Use EMMS

These guidelines help you determine when to use **EMMS**:

- *If next instruction is FP*—Use `_mm_empty()` after an MMX instruction if the next instruction is an FP instruction; for example, before doing calculations on floats, doubles or long doubles.
- *Don't empty when already empty*—If the next instruction uses an MMX register, `_mm_empty()` incurs an operation with no benefit (no-op).
- *Group Instructions*—Use different functions for regions that use FP instructions and those that use MMX instructions. This eliminates needing an EMMS instruction within the body of a critical loop.
- *Runtime initialization*—Use `_mm_empty()` during runtime initialization of `__m64` and FP data types. This ensures resetting the register between data type transitions. See usage coding [Example 13-1](#).

Example 13-1 Correct EMMS Usage In Initialization Code

Incorrect Usage	Correct Usage
<code>__m64 x = _m_padd(y, z);</code>	<code>__m64 x = _m_padd(y, z);</code>
<code>float f = init();</code>	<code>float f = (_mm_empty(), init());</code>

Further, you must be aware of all the situations when your code generates an MMX instruction with the Intel C/C++ compiler:

- when using an MMX intrinsic
- when using the Streaming SIMD Extensions (for those intrinsics that use MMX data)
- when using an MMX instruction through inline assembly
- when referencing an `__m64` data type variable

For more documentation on EMMS, visit the following web site:

<http://developer.intel.com>

MMX Technology Intrinsic Groups

The MMX technology intrinsics are grouped into the following sets:

- General Support Intrinsics
- Packed Arithmetic Intrinsics
- Shift Intrinsics
- Logical Intrinsics
- Compare Intrinsics

The syntax is defined for each intrinsic before the definition as follows:

```
data_type intrinsic_name (parameters) INST
```

Where,

data_type Is the return data type, which can be either `void`, `int`, or `__m64`. Only the `__mm_empty` intrinsic returns `void`, and only `__mm_cvtsi64_si32` returns an `int`. All other MMX technology intrinsics use the `__m64` data type.

intrinsic_name Is the name of the intrinsic, which behaves like a function that you can use in your C/C++ code instead of inlining the actual instruction. A table in each section provides the intrinsic names, alternate names, and the corresponding instruction. The names listed in the descriptions have a new mnemonic spelling that more easily indicates the functionality. The alternate names are the original intrinsic names, which are provided in case they are the ones that you have previously used.

INST Is the name of the assembly instruction that is used by the intrinsic. Do not use this instruction as part of the syntax in your code. The instruction is placed here to indicate the instruction to which that particular intrinsic corresponds.

General Support Intrinsics

The general support intrinsics are listed in [Table 13-1](#), and are followed by a description of each intrinsic.

Table 13-1 **General Support Intrinsics**

Intrinsic Name	Alternate Name	Corresponding Instruction
<code>_mm_empty</code>	<code>_m_empty</code>	EMMS
<code>_mm_cvtsi32_si64</code>	<code>_m_from_int</code>	MOVD
<code>_mm_cvtsi64_si32</code>	<code>_m_to_int</code>	MOVD
<code>_mm_packs_pi16</code>	<code>_m_packsswb</code>	PACKSSWB
<code>_mm_packs_pi32</code>	<code>_m_packssdw</code>	PACKSSDW
<code>_mm_packs_pul6</code>	<code>_m_packuswb</code>	PACKUSWB
<code>_mm_unpackhi_pi8</code>	<code>_m_punpckhbw</code>	PUNPCKHBW
<code>_mm_unpackhi_pi16</code>	<code>_m_punpckhwd</code>	PUNPCKHWD
<code>_mm_unpackhi_pi32</code>	<code>_m_punpckhdq</code>	PUNPCKHDQ
<code>_mm_unpacklo_pi8</code>	<code>_m_punpcklbw</code>	PUNPCKLBW
<code>_mm_unpacklo_pi16</code>	<code>_m_punpcklwd</code>	PUNPCKLWD
<code>_mm_unpacklo_pi32</code>	<code>_m_punpckldq</code>	PUNPCKLDQ

```
void _mm_empty (void) EMMS
```

Empty the multimedia state.

See [“The EMMS Instruction: Why You Need it and When to Use it”](#).

```
__m64 _mm_cvtsi32_si64 (int i) MOVD
```

Convert the integer object *i* to a 64-bit `__m64` object. The integer value is zero-extended to 64 bits.

```
int _mm_cvtsi64_si32 (__m64 m) MOVD
```

Convert the lower 32 bits of the `__m64` object *m* to an integer.

```
__m64 _mm_packs_pi16 (__m64 m1, __m64 m2) PACKSSWB
```

Pack the four 16-bit values from *m1* into the lower four 8-bit values of the result with signed saturation, and pack the four 16-bit values from *m2* into the upper four 8-bit values of the result with signed saturation.

```
__m64 _mm_packs_pi32 (__m64 m1, __m64 m2)    PACKSSDW
```

Pack the two 32-bit values from *m1* into the lower two 16-bit values of the result with signed saturation, and pack the two 32-bit values from *m2* into the upper two 16-bit values of the result with signed saturation.

```
__m64 _mm_packs_pu16 (__m64 m1, __m64 m2)    PACKUSWB
```

Pack the four 16-bit values from *m1* into the lower four 8-bit values of the result with unsigned saturation, and pack the four 16-bit values from *m2* into the upper four 8-bit values of the result with unsigned saturation.

```
__m64 _mm_unpackhi_pi8 (__m64 m1, __m64 m2)  PUNPCKHBW
```

Interleave the four 8-bit values from the high half of *m1* with the four values from the high half of *m2* and take the least significant element from *m1*.

```
__m64 _mm_unpackhi_pi16 (__m64 m1, __m64 m2) PUNPCKHWD
```

Interleave the two 16-bit values from the high half of *m1* with the two values from the high half of *m2* and take the least significant element from *m1*.

```
__m64 _mm_unpackhi_pi32 (__m64 m1, __m64 m2) PUNPCKHDQ
```

Interleave the 32-bit value from the high half of *m1* with the 32-bit value from the high half of *m2* and take the least significant element from *m1*.

```
__m64 _mm_unpacklo_pi8 (__m64 m1, __m64 m2)  PUNPCKLBW
```

Interleave the four 8-bit values from the low half of *m1* with the four values from the low half of *m2* and take the least significant element from *m1*.

```
__m64 _mm_unpacklo_pi16 (__m64 m1, __m64 m2) PUNPCKLWD
```

Interleave the two 16-bit values from the low half of *m1* with the two values from the low half of *m2* and take the least significant element from *m1*.

```
__m64 _mm_unpacklo_pi32 (__m64 m1, __m64 m2) PUNPCKLDQ
```

Interleave the 32-bit value from the low half of *m1* with the 32-bit value from the low half of *m2* and take the least significant element from *m1*.

Packed Arithmetic Intrinsics

The packed arithmetic intrinsics are listed in [Table 13-2](#), and are followed by a description of each intrinsic.

Table 13-2 Packed Arithmetic Intrinsics

Intrinsic Name	Alternate Name	Corresponding Instruction
<code>_mm_add_pi8</code>	<code>_m_paddb</code>	PADDB
<code>_mm_add_pi16</code>	<code>_m_paddw</code>	PADDW
<code>_mm_add_pi32</code>	<code>_m_paddd</code>	PADDQ
<code>_mm_adds_pi8</code>	<code>_m_paddsb</code>	PADDSSB
<code>_mm_adds_pi16</code>	<code>_m_paddsw</code>	PADDSSW
<code>_mm_adds_pu8</code>	<code>_m_paddusb</code>	PADDUSB
<code>_mm_adds_pu16</code>	<code>_m_paddusw</code>	PADDUSW
<code>_mm_sub_pi8</code>	<code>_m_psubb</code>	PSUBB
<code>_mm_sub_pi16</code>	<code>_m_psubw</code>	PSUBW
<code>_mm_sub_pi32</code>	<code>_m_psubd</code>	PSUBQ
<code>_mm_subs_pi8</code>	<code>_m_psubsb</code>	PSUBSSB
<code>_mm_subs_pi16</code>	<code>_m_psubsw</code>	PSUBSSW
<code>_mm_subs_pu8</code>	<code>_m_psubusb</code>	PSUBUSB
<code>_mm_subs_pu16</code>	<code>_m_psubusw</code>	PSUBUSW
<code>_mm_madd_pi16</code>	<code>_m_pmaddwd</code>	PMADDWD
<code>_mm_mulhi_pi16</code>	<code>_m_pmulhw</code>	PMULHW
<code>_mm_mullo_pi16</code>	<code>_m_pmullw</code>	PMULLW

<code>__m64 _mm_add_pi8 (__m64 m1, __m64 m2)</code>	PADDB
Add the eight 8-bit values in <i>m1</i> to the eight 8-bit values in <i>m2</i> .	
<code>__m64 _mm_add_pi16 (__m64 m1, __m64 m2)</code>	PADDW

Add the four 16-bit values in *m1* to the four 16-bit values in *m2*.

```
__m64 _mm_add_pi32 (__m64 m1, __m64 m2)          PADDQ
```

Add the two 32-bit values in *m1* to the two 32-bit values in *m2*.

```
__m64 _mm_adds_pi8 (__m64 m1, __m64 m2)          PADDSB
```

Add the eight signed 8-bit values in *m1* to the eight signed 8-bit values in *m2* and saturate.

```
__m64 _mm_adds_pi16 (__m64 m1, __m64 m2)         PADDSW
```

Add the four signed 16-bit values in *m1* to the four signed 16-bit values in *m2* and saturate.

```
__m64 _mm_adds_pu8 (__m64 m1, __m64 m2)          PADDUSB
```

Add the eight unsigned 8-bit values in *m1* to the eight unsigned 8-bit values in *m2* and saturate.

```
__m64 _mm_adds_pu16 (__m64 m1, __m64 m2)         PADDUSW
```

Add the four unsigned 16-bit values in *m1* to the four unsigned 16-bit values in *m2* and saturate.

```
__m64 _mm_sub_pi8 (__m64 m1, __m64 m2)           PSUBB
```

Subtract the eight 8-bit values in *m2* from the eight 8-bit values in *m1*.

```
__m64 _mm_sub_pi16 (__m64 m1, __m64 m2)          PSUBW
```

Subtract the four 16-bit values in *m2* from the four 16-bit values in *m1*.

```
__m64 _mm_sub_pi32 (__m64 m1, __m64 m2)          PSUBD
```

Subtract the two 32-bit values in *m2* from the two 32-bit values in *m1*.

```
__m64 _mm_subs_pi8 (__m64 m1, __m64 m2)          PSUBSB
```

Subtract the eight signed 8-bit values in *m2* from the eight signed 8-bit values in *m1* and saturate.

```
__m64 _mm_subs_pi16 (__m64 m1, __m64 m2)         PSUBSW
```

Subtract the four signed 16-bit values in *m2* from the four signed 16-bit values in *m1* and saturate.

`__m64 _mm_subs_pu8 (__m64 m1, __m64 m2)` PSUBUSB

Subtract the eight unsigned 8-bit values in *m2* from the eight unsigned 8-bit values in *m1* and saturate.

`__m64 _mm_subs_pu16 (__m64 m1, __m64 m2)` PSUBUSW

Subtract the four unsigned 16-bit values in *m2* from the four unsigned 16-bit values in *m1* and saturate.

`__m64 _mm_madd_pi16 (__m64 m1, __m64 m2)` PMADDWD

Multiply four 16-bit values in *m1* by four 16-bit values in *m2* producing four 32-bit intermediate results, which are then summed by pairs to produce two 32-bit results.

`__m64 _mm_mulhi_pi16 (__m64 m1, __m64 m2)` PMULHW

Multiply four signed 16-bit values in *m1* by four signed 16-bit values in *m2* and produce the high 16 bits of the four results.

`__m64 _mm_mullo_pi16 (__m64 m1, __m64 m2)` PMULLW

Multiply four 16-bit values in *m1* by four 16-bit values in *m2* and produce the low 16 bits of the four results.

Shift Intrinsics

The shift intrinsics are listed in [Table](#), and are followed by a description of each intrinsic.

Table 13-3 **Shift Intrinsics**

Intrinsic Name	Alternate Name	Corresponding Instruction
<code>_mm_sll_pi16</code>	<code>_m_psllw</code>	PSLLW
<code>_mm_slli_pi16</code>	<code>_m_psllwi</code>	PSLLW
<code>_mm_sll_pi32</code>	<code>_m_psll_d</code>	PSLLD
<code>_mm_slli_pi32</code>	<code>_m_psll_d_i</code>	PSLLD

Table 13-3 Shift Intrinsics (Continued)

Intrinsic Name	Alternate Name	Corresponding Instruction
<code>_mm_sll_si64</code>	<code>_m_psllq</code>	PSLLQ
<code>_mm_slli_si64</code>	<code>_m_psllqi</code>	PSLLQ
<code>_mm_sra_pi16</code>	<code>_m_psraw</code>	PSRAW
<code>_mm_srai_pi16</code>	<code>_m_psrawi</code>	PSRAW
<code>_mm_sra_pi32</code>	<code>_m_psrad</code>	PSRAD
<code>_mm_srai_pi32</code>	<code>_m_psradi</code>	PSRADI
<code>_mm_srl_pi16</code>	<code>_m_psrlw</code>	PSRLW
<code>_mm_srli_pi16</code>	<code>_m_psrlwi</code>	PSRLW
<code>_mm_srl_pi32</code>	<code>_m_psrld</code>	PSRLD
<code>_mm_srli_pi32</code>	<code>_m_psrldi</code>	PSRLD
<code>_mm_srl_si64</code>	<code>_m_psrlq</code>	PSRLQ
<code>_mm_srli_si64</code>	<code>_m_psrlqi</code>	PSRLQ

`__m64 _mm_sll_pi16 (__m64 m, __m64 count)` PSLLW

Shift four 16-bit values in *m* left the amount specified by *count* while shifting in zeros.

`__m64 _mm_slli_pi16 (__m64 m, int count)` PSLLW

Shift four 16-bit values in *m* left the amount specified by *count* while shifting in zeros. For the best performance, *count* should be a constant.

`__m64 _mm_sll_pi32 (__m64 m, __m64 count)` PSLLD

Shift two 32-bit values in *m* left the amount specified by *count* while shifting in zeros.

`__m64 _mm_slli_pi32 (__m64 m, int count)` PSLLD

Shift two 32-bit values in *m* left the amount specified by *count* while shifting in zeros. For the best performance, *count* should be a constant.

`__m64 _mm_sll_si64 (__m64 m, __m64 count)` PSLLQ

Shift the 64-bit value in *m* left the amount specified by *count* while shifting in zeros.

`__m64 _mm_slli_si64 (__m64 m, int count)` PSLAQ

Shift the 64-bit value in *m* left the amount specified by *count* while shifting in zeros. For the best performance, *count* should be a constant.

`__m64 _mm_sra_pi16 (__m64 m, __m64 count)` PSRAW

Shift four 16-bit values in *m* right the amount specified by *count* while shifting in the sign bit.

`__m64 _mm_srai_pi16 (__m64 m, int count)` PSRAW

Shift four 16-bit values in *m* right the amount specified by *count* while shifting in the sign bit. For the best performance, *count* should be a constant.

`__m64 _mm_sra_pi32 (__m64 m, __m64 count)` PSRAD

Shift two 32-bit values in *m* right the amount specified by *count* while shifting in the sign bit.

`__m64 _mm_srai_pi32 (__m64 m, int count)` PSRAI

Shift two 32-bit values in *m* right the amount specified by *count* while shifting in the sign bit. For the best performance, *count* should be a constant.

`__m64 _mm_srl_pi16 (__m64 m, __m64 count)` PSRLW

Shift four 16-bit values in *m* right the amount specified by *count* while shifting in zeros.

`__m64 _mm_srli_pi16 (__m64 m, int count)` PSRLW

Shift four 16-bit values in *m* right the amount specified by *count* while shifting in zeros. For the best performance, *count* should be a constant.

`__m64 _mm_srl_pi32 (__m64 m, __m64 count)` PSRLD

Shift two 32-bit values in *m* right the amount specified by *count* while shifting in zeros.

```
__m64 _mm_srli_pi32 (__m64 m, int count) PSRLD
```

Shift two 32-bit values in *m* right the amount specified by *count* while shifting in zeros. For the best performance, *count* should be a constant.

```
__m64 _mm_srl_si64 (__m64 m, __m64 count) PSRLQ
```

Shift the 64-bit value in *m* right the amount specified by *count* while shifting in zeros.

```
__m64 _mm_srli_si64 (__m64 m, int count) PSRLQ
```

Shift the 64-bit value in *m* right the amount specified by *count* while shifting in zeros. For the best performance, *count* should be a constant.

Logical Intrinsics

The logical intrinsics are listed in [Table 13-4](#), and are followed by a description of each intrinsic

Table 13-4 Logical Intrinsics

Intrinsic Name	Alternate Name	Corresponding Instruction
<code>_mm_and_si64</code>	<code>_m_pand</code>	PAND
<code>_mm_andnot_si64</code>	<code>_m_pandn</code>	PANDN
<code>_mm_or_si64</code>	<code>_m_por</code>	POR
<code>_mm_xor_si64</code>	<code>_m_pxor</code>	PXOR

```
__m64 _mm_and_si64 (__m64 m1, __m64 m2) PAND
```

Perform a bitwise **AND** of the 64-bit value in *m1* with the 64-bit value in *m2*.

```
__m64 _mm_andnot_si64 (__m64 m1, __m64 m2) PANDN
```

Perform a logical **NOT** on the 64-bit value in *m1* and use the result in a bitwise **AND** with the 64-bit value in *m2*.

```
__m64 _mm_or_si64 (__m64 m1, __m64 m2) POR
```

Perform a bitwise **OR** of the 64-bit value in *m1* with the 64-bit value in *m2*.

```
__m64 _mm_xor_si64 (__m64 m1, __m64 m2)          PXOR
```

Perform a bitwise **XOR** of the 64-bit value in *m1* with the 64-bit value in *m2*.

Compare Intrinsics

The compare intrinsics are listed in [Table 13-4](#), and are followed by a description of each intrinsic.

Table 13-5 Compare Intrinsics

Intrinsic Name	Alternate Name	Corresponding Instruction
<code>_mm_cmpeq_pi8</code>	<code>_m_pcmpeqb</code>	PCMPEQB
<code>_mm_cmpeq_pi16</code>	<code>_m_pcmpeqw</code>	PCMPEQW
<code>_mm_cmpeq_pi32</code>	<code>_m_pcmpeqd</code>	PCMPEQD
<code>_mm_cmpgt_pi8</code>	<code>_m_pcmpgtb</code>	PCMPGTB
<code>_mm_cmpgt_pi16</code>	<code>_m_pcmpgtw</code>	PCMPGTW
<code>_mm_cmpgt_pi32</code>	<code>_m_pcmpgtd</code>	PCMPGTD

```
__m64 _mm_cmpeq_pi8 (__m64 m1, __m64 m2)          PCMPEQB
```

If the respective 8-bit values in *m1* are equal to the respective 8-bit values in *m2* set the respective 8-bit resulting values to all ones, otherwise set them to all zeros.

```
__m64 _mm_cmpeq_pi16 (__m64 m1, __m64 m2)         PCMPEQW
```

If the respective 16-bit values in *m1* are equal to the respective 16-bit values in *m2* set the respective 16-bit resulting values to all ones, otherwise set them to all zeros.

```
__m64 _mm_cmpeq_pi32 (__m64 m1, __m64 m2)         PCMPEQD
```

If the respective 32-bit values in *m1* are equal to the respective 32-bit values in *m2* set the respective 32-bit resulting values to all ones, otherwise set them to all zeros.

```
__m64 _mm_cmpgt_pi8 (__m64 m1, __m64 m2)          PCMPGTB
```

If the respective 8-bit values in *m1* are greater than the respective 8-bit values in *m2* set the respective 8-bit resulting values to all ones, otherwise set them to all zeros.

```
__m64 _mm_cmpgt_pi16 (__m64 m1, __m64 m2)      PCMPGTW
```

If the respective 16-bit values in *m1* are greater than the respective 16-bit values in *m2* set the respective 16-bit resulting values to all ones, otherwise set them to all zeros.

```
__m64 _mm_cmpgt_pi32 (__m64 m1, __m64__m64 m2)  PCMPGTD
```

If the respective 32-bit values in *m1* are greater than the respective 32-bit values in *m2* set the respective 32-bit resulting values to all ones, otherwise set them all to zeros.

Processor-Dispatching Support

With `__declspec(cpu_specific)` and `__declspec(cpu_dispatch)`, you can direct the compiler to generate code that automatically determines the processor on which it is running, and selects the appropriate implementation of a function. This enables you, for example, to write code that takes advantage of MMX instructions when executing on a processor that has them, but that also executes correctly on older processors that do not have the MMX instructions. The syntax of these extended-attributes is as follows:

```
cpu_specific(cpuid)
cpu_dispatch(cpuid-list)
```

where *cpuid* is one of:

```
generic
pentium
pentium_pro
pentium_mmx
pentium_ii
pentium_iii
pentium_iii_no_xmm_regs
```

and *cpuid-list* is one of:

cpuid
cpuid-list , *cpuid*

The names of the *cpuid* are not case sensitive. You can specify these attributes only for function definitions. The body of a function declared with `__declspec(cpu_dispatch)` must be empty, and is referred to as a stub.

If a function *f* is defined as `__declspec(cpu_specific(p))`, then a *cpu_dispatch* stub must appear for *f* somewhere in the program, and *p* must be in the *cpuid-list* of that stub; otherwise, that *cpu_specific* definition can never be called and no error will be reported for this condition.

If the *cpu_dispatch* stub for a function *f* contains the *cpuid p*, then a *cpu_specific* definition of *f* with *cpuid p* must appear somewhere in the program; otherwise an unresolved external error is reported. A *cpu_specific* function definition need not appear in the same translation unit as the corresponding *cpu_dispatch* stub, unless the *cpu_specific* function is declared *static*. The *inline* attribute is disabled for all *cpu_specific* and *cpu_dispatch* functions.

When a *cpu_dispatch* stub is compiled, its body is filled in with code that determines the processor on which the program is running, then dispatches to the “best” *cpu_specific* implementation available (as defined by the *cpuid-list*) that will run on that processor. When a *cpu_specific* function is compiled, optimizations and code generation options appropriate to the specified processor are used regardless of command-line option settings.

Here is an example of how these features can be used:

```
#include <mmintrin.h>

/* array_sum(r, a, b, l) adds two arrays of unsigned
short (a and b), of length l, and stores the result in r.
*/

__declspec(cpu_specific(Pentium))
void array_sum(unsigned short *result,
unsigned short const *a,
```

```

    unsigned short const *b,
    size_t length)

{
    /* The implementation specific to the Pentium processor
    uses no special architectural features. */

    for (; length > 0; length--)
        *result++ = *a++ + *b++;
}

__declspec(cpu_specific(Pentium_MMX))
void array_sum(unsigned short *result,
unsigned short const *a,
unsigned short const *b,
size_t length)
{
    /* The implementation for a Pentium processor with MMX
    technology uses an MMX instruction intrinsic to add
    four elements at a time, to reduce the number of loop
    iterations by 3/4. */

    __m64 *mmx_result = (__m64 *)result;
    __m64 const *mmx_a = (__m64 const *)a;
    __m64 const *mmx_b = (__m64 const *)b;
    for (; length > 3; length -= 4)
        *mmx_result++ = _m_paddw(*mmx_a++, *mmx_b++);

    /* If the size of all arrays passed to this routine is
    known to be a multiple of four, the following code
    (which takes care of excess elements) is not necessary.
    */

    result = (unsigned short *)mmx_result;
    a = (unsigned short const *)mmx_a;
    b = (unsigned short const *)mmx_b;
    for (; length > 0; length--)
        *result++ = *a++ + *b++;
}

```

```
}

__declspec(cpu_dispatch(Pentium, Pentium_MMX))
void array_sum(unsigned short *result,
unsigned short const *a,
unsigned short const *b,
size_t length)
{
/* An empty function body, which informs the compiler
that it should generate a dispatch function for the
CPU-specific implementations listed in the cpu_dispatch
clause. */
}
```

Streaming SIMD Extensions Intrinsics

This section describes the C/C++ language-level features supporting the Streaming SIMD Extensions in the Intel C/C++ Compiler. The section explains the following features of the intrinsics:

- the intrinsics API
- `__m128` datatype
- data alignment support
- assembly language support

This section provides the complete list of intrinsics for the Streaming SIMD Extensions.

The Intrinsics API

The intrinsics for the Streaming SIMD Extensions allow you to access the functionality provided by the new instructions without actually using assembly language. Just like the intrinsics for the integer MMX instructions, for each computational and data manipulation instruction in the new instruction set, there is a corresponding C intrinsic that implements it directly. A new C data type, representing the 128-bit registers, is used as

the operand to these intrinsic functions. The intrinsics allow you to specify the underlying implementation (instruction selection) of an algorithm yet leave instruction scheduling and register allocation to the compiler.

The `__m128` Data Type

The `__m128` data type is used to represent the contents of an `xmm` register, which is either four packed single-precision floating-point values or one scalar single-precision number. The `__m128` data type is not a basic ANSI C data type and therefore some restrictions are placed on its usage:

- Use `__m128` only on the left-hand side of an assignment, as a return value, or as a parameter. Do not use it in other arithmetic expressions such as "+" and ">>".
- Do not initialize `__m128` with literals; there is no way to express 128-bit constants.
- Use `__m128` objects in aggregates, such as unions (for example, to access the float elements) and structures. The address of an `__m128` object may be taken.
- Use `__m128` data only with the intrinsics described in this user's guide.

The compiler aligns `__m128` local data to 16B boundaries on the stack. Global `__m128` data is also 16B-aligned. (To align float arrays, you can use the alignment `declspec` described in the following section.) Because the new instruction set treats the Streaming SIMD Extensions registers in the same way whether you are using packed or scalar data, there is no `__m32` datatype to represent scalar data as you might expect. For scalar operations, you should use the `__m128` objects and the “scalar” forms of the intrinsics; the compiler and the processor implement these operations with 32-bit memory references.

Streaming SIMD Extensions Intrinsic Conventions

The list of intrinsics for the Streaming SIMD Extensions uses a number of conventions described in the following paragraphs.

The suffixes `ps` and `ss` are used to denote “packed single” and “scalar single” precision operations. The packed floats are represented in right-to-left order, with the lowest word (right-most) being used for scalar operations: [`z`, `y`, `x`, `w`]. To explain how memory storage reflects this, consider the following example. The operation

```
float a[4] = { 1.0, 2.0, 3.0, 4.0 };
```

```
__m128 t = _mm_load_ps(a);
```

produces the same result as follows:

```
__m128 t = _mm_set_ps(4.0, 3.0, 2.0, 1.0);
```

In other words,

```
t = [ 4.0, 3.0, 2.0, 1.0 ]
```

where the “scalar” element is 1.0.

Floating Point Intrinsics

The following sections list the floating-point intrinsics broken into groups by the nature of the operation. Each intrinsic entry has an informal pseudo-code and it is followed with a corresponding instruction name in upper case letters; for example, **ADDSS** is the name of the first instruction listed in this section, which corresponds to the intrinsic for

```
__m128 _mm_add_ss(__m128 a, __m128 b).
```

The variable **r** is generally used for the intrinsic’s return value. A number appended to a variable name indicates the element of a packed object. For example, **r0** is the lowest word of **r**. Some intrinsics are “composites” because they require more than one instruction to implement them.

You should be familiar with the hardware features provided by the Streaming SIMD Extensions when writing programs with the intrinsics. The following are three important issues to keep in mind:

- Certain intrinsics, such as `_mm_loadr_ps` and `_mm_cmpgt_ss`, are not directly supported by the instruction set. While these intrinsics are convenient programming aids, be mindful that they might consist of more than one machine-language instruction.
- Floating-point data loaded or stored as `__m128` objects must be generally 16-byte-aligned.
- Some intrinsics require that their argument be immediates, that is, constant integers (literals), due to the nature of the instruction.

- The result of arithmetic operations acting on two NaN (Not a Number) arguments is undefined. Therefore, FP operations using NaN arguments will not match the expected behavior of the corresponding assembly instructions.

Arithmetic Operations

`__m128 _mm_add_ss(__m128 a, __m128 b)` ADDSS

Adds the lower SP FP (single-precision, floating-point) values of *a* and *b*; the upper 3 SP FP values are passed through from *a*.

`r0 := a0 + b0`

`r1 := a1 ; r2 := a2 ; r3 := a3`

`__m128 _mm_add_ps(__m128 a, __m128 b)` ADDPS

Adds the four SP FP values of *a* and *b*.

`r0 := a0 + b0`

`r1 := a1 + b1`

`r2 := a2 + b2`

`r3 := a3 + b3`

`__m128 _mm_sub_ss(__m128 a, __m128 b)` SUBSS

Subtracts the lower SP FP values of *a* and *b*. The upper 3 SP FP values are passed through from *a*.

`r0 := a0 - b0`

`r1 := a1 ; r2 := a2 ; r3 := a3`

`__m128 _mm_sub_ps(__m128 a, __m128 b)` SUBPS

Subtracts the four SP FP values of *a* and *b*.

`r0 := a0 - b0`

`r1 := a1 - b1`

`r2 := a2 - b2`

`r3 := a3 - b3`

`__m128 _mm_mul_ss(__m128 a, __m128 b)` MULSS

Multiplies the lower SP FP values of *a* and *b*; the upper 3 SP FP values are passed through from *a*.

```
r0 := a0 * b0  
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_mul_ps(__m128 a, __m128 b) MULPS
```

Multiplies the four SP FP values of *a* and *b*.

```
r0 := a0 * b0  
r1 := a1 * b1  
r2 := a2 * b2  
r3 := a3 * b3
```

```
__m128 _mm_div_ss(__m128 a, __m128 b) DIVSS
```

Divides the lower SP FP values of *a* and *b*; the upper 3 SP FP values are passed through from *a*.

```
r0 := a0 / b0  
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_div_ps(__m128 a, __m128 b) DIVPS
```

Divides the four SP FP values of *a* and *b*.

```
r0 := a0 / b0  
r1 := a1 / b1  
r2 := a2 / b2  
r3 := a3 / b3
```

```
__m128 _mm_sqrt_ss(__m128 a) SQRTSS
```

Computes the square root of the lower SP FP value of *a*; the upper 3 SP FP values are passed through.

```
r0 := sqrt(a0)  
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_sqrt_ps(__m128 a) SQRTPS
```

Computes the square roots of the four SP FP values of *a*.

```
r0 := sqrt(a0)
```

```
r1 := sqrt(a1)
```

```
r2 := sqrt(a2)
```

```
r3 := sqrt(a3)
```

```
__m128 _mm_rcp_ss(__m128 a)
```

RCPSS

Computes the approximation of the reciprocal of the lower SP FP value of **a**; the upper 3 SP FP values are passed through.

```
r0 := recip(a0)
```

```
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_rcp_ps(__m128 a)
```

RCPPS

Computes the approximations of reciprocals of the four SP FP values of **a**.

```
r0 := recip(a0)
```

```
r1 := recip(a1)
```

```
r2 := recip(a2)
```

```
r3 := recip(a3)
```

```
__m128 _mm_rsqrt_ss(__m128 a)
```

RSQRTSS

Computes the approximation of the reciprocal of the square root of the lower SP FP value of **a**; the upper 3 SP FP values are passed through.

```
r0 := recip(sqrt(a0))
```

```
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_rsqrt_ps(__m128 a)
```

RSQRTPS

Computes the approximations of the reciprocals of the square roots of the four SP FP values of **a**.

```
r0 := recip(sqrt(a0))
```

```
r1 := recip(sqrt(a1))
```

```
r2 := recip(sqrt(a2))
```

```
r3 := recip(sqrt(a3))
```

```
__m128 _mm_min_ss(__m128 a, __m128 b)
```

MINSS

Computes the minimum of the lower SP FP values of **a** and **b**; the upper 3 SP FP values are passed through from **a**.

```
r0 := min(a0, b0)
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_min_ps(__m128 a, __m128 b) MINPS
```

Computes the minimums of the four SP FP values of *a* and *b*.

```
r0 := min(a0, b0)
r1 := min(a1, b1)
r2 := min(a2, b2)
r3 := min(a3, b3)
```

```
__m128 _mm_max_ss(__m128 a, __m128 b) MAXSS
```

Computes the maximum of the lower SP FP values of *a* and *b*; the upper 3 SP FP values are passed through from *a*.

```
r0 := max(a0, b0)
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_max_ps(__m128 a, __m128 b) MAXPS
```

Computes the maximums of the four SP FP values of *a* and *b*.

```
r0 := max(a0, b0)
r1 := max(a1, b1)
r2 := max(a2, b2)
r3 := max(a3, b3)
```

Logical Operations

```
__m128 _mm_and_ps(__m128 a, __m128 b) ANDPS
```

Computes the bitwise And of the four SP FP values of *a* and *b*.

```
r0 := a0 & b0
r1 := a1 & b1
r2 := a2 & b2
r3 := a3 & b3
```

```
__m128 _mm_andnot_ps(__m128 a, __m128 b) ANDNPS
```

Computes the bitwise AND-NOT of the four SP FP values of *a* and *b*.

```
r0 := ~a0 & b0
r1 := ~a1 & b1
r2 := ~a2 & b2
r3 := ~a3 & b3
```

```
__m128 _mm_or_ps(__m128 a, __m128 b) ORPS
```

Computes the bitwise OR of the four SP FP values of *a* and *b*.

```
r0 := a0 | b0
r1 := a1 | b1
r2 := a2 | b2
r3 := a3 | b3
```

```
__m128 _mm_xor_ps(__m128 a, __m128 b) XORPS
```

Computes bitwise EXOR (exclusive-or) of the four SP FP values of *a* and *b*.

```
r0 := a0 ^ b0
r1 := a1 ^ b1
r2 := a2 ^ b2
r3 := a3 ^ b3
```

Comparisons

Each comparison intrinsic performs a comparison of *a* and *b*. For the packed form, the four SP FP values of *a* and *b* are compared, and a 128-bit mask is returned. For the scalar form, the lower SP FP values of *a* and *b* are compared, and a 32-bit mask is returned; the upper three SP FP values are passed through from *a*. The mask is set to `0xffffffff` for each element where the comparison is true and `0x0` where the comparison is false.

The superscript 'r' on the instruction indicates operands are reversed in the instruction implementation.

```
__m128 _mm_cmpeq_ss(__m128 a, __m128 b) CMPEQSS
```

Compare for equality.

```
r0 := (a0 == b0) ? 0xffffffff : 0x0
```

```
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cmpeq_ps(__m128 a, __m128 b) CMPEQPS
```

Compare for equality.

```
r0 := (a0 == b0) ? 0xffffffff : 0x0
```

```
r1 := (a1 == b1) ? 0xffffffff : 0x0
```

```
r2 := (a2 == b2) ? 0xffffffff : 0x0
```

```
r3 := (a3 == b3) ? 0xffffffff : 0x0
```

```
__m128 _mm_cmplt_ss(__m128 a, __m128 b) CMPLTSS
```

Compare for less-than.

```
r0 := (a0 < b0) ? 0xffffffff : 0x0
```

```
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cmplt_ps(__m128 a, __m128 b) CMPLTPS
```

Compare for less-than.

```
r0 := (a0 < b0) ? 0xffffffff : 0x0
```

```
r1 := (a1 < b1) ? 0xffffffff : 0x0
```

```
r2 := (a2 < b2) ? 0xffffffff : 0x0
```

```
r3 := (a3 < b3) ? 0xffffffff : 0x0
```

```
__m128 _mm_cmple_ss(__m128 a, __m128 b) CMPLESS
```

Compare for less-than-or-equal.

```
r0 := (a0 <= b0) ? 0xffffffff : 0x0
```

```
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cmple_ps(__m128 a, __m128 b) CMPLEPS
```

Compare for less-than-or-equal.

```
r0 := (a0 <= b0) ? 0xffffffff : 0x0
```

```
r1 := (a1 <= b1) ? 0xffffffff : 0x0
```

```
r2 := (a2 <= b2) ? 0xffffffff : 0x0
```

```
r3 := (a3 <= b3) ? 0xffffffff : 0x0
```



```
__m128 _mm_cmpgt_ss(__m128 a, __m128 b) CMPLTSSr
```

Compare for greater-than.

```
r0 := (a0 > b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cmpgt_ps(__m128 a, __m128 b) CMPLTPSr
```

Compare for greater-than.

```
r0 := (a0 > b0) ? 0xffffffff : 0x0
r1 := (a1 > b1) ? 0xffffffff : 0x0
r2 := (a2 > b2) ? 0xffffffff : 0x0
r3 := (a3 > b3) ? 0xffffffff : 0x0
```

```
__m128 _mm_cmpge_ss(__m128 a, __m128 b) CMPLESSr
```

Compare for greater-than-or-equal.

```
r0 := (a0 >= b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cmpge_ps(__m128 a, __m128 b) CMPLEPSr
```

Compare for greater-than-or-equal.

```
r0 := (a0 >= b0) ? 0xffffffff : 0x0
r1 := (a1 >= b1) ? 0xffffffff : 0x0
r2 := (a2 >= b2) ? 0xffffffff : 0x0
r3 := (a3 >= b3) ? 0xffffffff : 0x0
```

```
__m128 _mm_cmpneq_ss(__m128 a, __m128 b) CMPNEQSS
```

Compare for inequality.

```
r0 := (a0 != b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cmpneq_ps(__m128 a, __m128 b) CMPNEQPS
```

Compare for inequality.

```
r0 := (a0 != b0) ? 0xffffffff : 0x0
r1 := (a1 != b1) ? 0xffffffff : 0x0
```

```
r2 := (a2 != b2) ? 0xffffffff : 0x0
r3 := (a3 != b3) ? 0xffffffff : 0x0
```

```
__m128 _mm_cmpnlt_ss(__m128 a, __m128 b)          CMPNLTSS
```

Compare for not-less-than.

```
r0 := !(a0 < b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cmpnlt_ps(__m128 a, __m128 b)          CMPNLTPS
```

Compare for not-less-than.

```
r0 := !(a0 < b0) ? 0xffffffff : 0x0
r1 := !(a1 < b1) ? 0xffffffff : 0x0
r2 := !(a2 < b2) ? 0xffffffff : 0x0
r3 := !(a3 < b3) ? 0xffffffff : 0x0
```

```
__m128 _mm_cmpnle_ss(__m128 a, __m128 b)          CMPNLESS
```

Compare for not-less-than-or-equal.

```
r0 := !(a0 <= b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cmpnle_ps(__m128 a, __m128 b)          CMPNLEPS
```

Compare for not-less-than-or-equal.

```
r0 := !(a0 <= b0) ? 0xffffffff : 0x0
r1 := !(a1 <= b1) ? 0xffffffff : 0x0
r2 := !(a2 <= b2) ? 0xffffffff : 0x0
r3 := !(a3 <= b3) ? 0xffffffff : 0x0
```

```
__m128 _mm_cmpngt_ss(__m128 a, __m128 b)          CMPNLTSSr
```

Compare for not-greater-than.

```
r0 := !(a0 > b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cmpngt_ps(__m128 a, __m128 b) CMPNLTPSr
```

Compare for not-greater-than.

```
r0 := !(a0 > b0) ? 0xffffffff : 0x0
r1 := !(a1 > b1) ? 0xffffffff : 0x0
r2 := !(a2 > b2) ? 0xffffffff : 0x0
r3 := !(a3 > b3) ? 0xffffffff : 0x0
```

```
__m128 _mm_cmpnge_ss(__m128 a, __m128 b) CMPNLESSr
```

Compare for not-greater-than-or-equal.

```
r0 := !(a0 >= b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cmpnge_ps(__m128 a, __m128 b) CMPNLEPSr
```

Compare for not-greater-than-or-equal.

```
r0 := !(a0 >= b0) ? 0xffffffff : 0x0
r1 := !(a1 >= b1) ? 0xffffffff : 0x0
r2 := !(a2 >= b2) ? 0xffffffff : 0x0
r3 := !(a3 >= b3) ? 0xffffffff : 0x0
```

```
__m128 _mm_cmpord_ss(__m128 a, __m128 b) CMPORDSS
```

Compare for ordered.

```
r0 := (a0 ord? b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cmpord_ps(__m128 a, __m128 b) CMPORDPS
```

Compare for ordered.

```
r0 := (a0 ord? b0) ? 0xffffffff : 0x0
r1 := (a1 ord? b1) ? 0xffffffff : 0x0
r2 := (a2 ord? b2) ? 0xffffffff : 0x0
r3 := (a3 ord? b3) ? 0xffffffff : 0x0
```

```
__m128 __mm_cmpunord_ss(__m128 a, __m128 b)    CMPUNORDSS
```

Compare for unordered.

```
r0 := (a0 unord? b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmpunord_ps(__m128 a, __m128 b)    CMPUNORDPS
```

Compare for unordered.

```
r0 := (a0 unord? b0) ? 0xffffffff : 0x0
r1 := (a1 unord? b1) ? 0xffffffff : 0x0
r2 := (a2 unord? b2) ? 0xffffffff : 0x0
r3 := (a3 unord? b3) ? 0xffffffff : 0x0
```

```
int __mm_comieq_ss (__m128 a, __m128 b)        COMISS
```

Compares the lower SP FP value of *a* and *b* for *a* equal to *b*. If *a* and *b* are equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 == b0) ? 0x1 : 0x0
```

```
int __mm_comilt_ss (__m128 a, __m128 b)        COMISS
```

Compares the lower SP FP value of *a* and *b* for *a* less than *b*. If *a* is less than *b*, 1 is returned. Otherwise 0 is returned.

```
r := (a0 < b0) ? 0x1 : 0x0
```

```
int __mm_comile_ss (__m128 a, __m128 b)        COMISS
```

Compares the lower SP FP value of *a* and *b* for *a* less than or equal to *b*. If *a* is less than or equal to *b*, 1 is returned. Otherwise 0 is returned.

```
r := (a0 <= b0) ? 0x1 : 0x0
```

```
int __mm_comigt_ss (__m128 a, __m128 b)        COMISS
```

Compares the lower SP FP value of *a* and *b* for *a* greater than *b*. If *a* is greater than *b* are equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 > b0) ? 0x1 : 0x0
```

```
int _mm_comige_ss (__m128 a, __m128 b) COMISS
```

Compares the lower SP FP value of *a* and *b* for *a* greater than or equal to *b*. If *a* is greater than or equal to *b*, 1 is returned. Otherwise 0 is returned.

```
r := (a0 >= b0) ? 0x1 : 0x0
```

```
int _mm_comineq_ss (__m128 a, __m128 b) COMISS
```

Compares the lower SP FP value of *a* and *b* for *a* not equal to *b*. If *a* and *b* are not equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 != b0) ? 0x1 : 0x0
```

```
int _mm_ucomieq_ss (__m128 a, __m128 b) UCOMISS
```

Compares the lower SP FP value of *a* and *b* for *a* equal to *b*. If *a* and *b* are equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 == b0) ? 0x1 : 0x0
```

```
int _mm_ucomilt_ss (__m128 a, __m128 b) UCOMISS
```

Compares the lower SP FP value of *a* and *b* for *a* less than *b*. If *a* is less than *b*, 1 is returned. Otherwise 0 is returned.

```
r := (a0 < b0) ? 0x1 : 0x0
```

```
int _mm_ucomile_ss (__m128 a, __m128 b) UCOMISS
```

Compares the lower SP FP value of *a* and *b* for *a* less than or equal to *b*. If *a* is less than or equal to *b*, 1 is returned. Otherwise 0 is returned.

```
r := (a0 <= b0) ? 0x1 : 0x0
```

```
int _mm_ucomigt_ss (__m128 a, __m128 b) UCOMISS
```

Compares the lower SP FP value of *a* and *b* for *a* greater than *b*. If *a* is greater than *b* are equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 > b0) ? 0x1 : 0x0
```

```
int _mm_ucomige_ss (__m128 a, __m128 b)          UCOMISS
```

Compares the lower SP FP value of *a* and *b* for *a* greater than or equal to *b*. If *a* is greater than or equal to *b*, 1 is returned. Otherwise 0 is returned.

```
r := (a0 >= b0) ? 0x1 : 0x0
```

```
int _mm_ucomineq_ss (__m128 a, __m128 b)         UCOMISS
```

Compares the lower SP FP value of *a* and *b* for *a* not equal to *b*. If *a* and *b* are not equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 != b0) ? 0x1 : 0x0
```

Conversion Operations

The conversions operations are listed in [Table 13-6](#), followed by a description of each intrinsic with the most recent mnemonic naming convention. The alternate name is provided in case you have used these intrinsics before.

Table 13-6 Conversion Operations

Intrinsic Name	Alternate Name	Corresponding Instruction
<code>_mm_cvtss_si32</code>	<code>_mm_cvt_ss2si</code>	CVTSS2SI
<code>_mm_cvtps_pi32</code>	<code>_mm_cvt_ps2pi</code>	CVTPS2PI
<code>_mm_cvttss_si32</code>	<code>_mm_cvtt_ss2si</code>	CVTTSS2SI
<code>_mm_cvttps_pi32</code>	<code>_mm_cvtt_ps2pi</code>	CVTTPS2PI
<code>_mm_cvtsi32_ss</code>	<code>_mm_cvt_si2ss</code>	CVTSI2SS
<code>_mm_cvtpi32_ps</code>	<code>_mm_cvt_pi2ps</code>	CVTTPS2PI

```
int _mm_cvtss_si32(__m128 a)                      CVTSS2SI
```

Convert the lower SP FP value of *a* to a 32-bit integer according to the current rounding mode.

```
r := (int)a0
```

```
__m64 _mm_cvtps_pi32(__m128 a)                    CVTPS2PI
```

Convert the two lower SP FP values of *a* to two 32-bit integers according to the current rounding mode, returning the integers in packed form.

```
r0 := (int)a0
```

```
r1 := (int)a1
```

```
int _mm_cvttsi_si32(__m128 a) CVTTSS2SI
```

Convert the lower SP FP value of *a* to a 32-bit integer with truncation.

```
r := (int)a0
```

```
__m64 _mm_cvttps_pi32(__m128 a) CVTTPS2PI
```

Convert the two lower SP FP values of *a* to two 32-bit integer with truncation, returning the integers in packed form.

```
r0 := (int)a0
```

```
r1 := (int)a1
```

```
__m128 _mm_cvtsi32_ss(__m128 a, int b) CVTSI2SS
```

Convert the 32-bit integer value *b* to an SP FP value; the upper three SP FP values are passed through from *a*.

```
r0 := (float)b
```

```
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cvtpi32_ps(__m128 a, __m64 b) CVTPi2PS
```

Convert the two 32-bit integer values in packed form in *b* to two SP FP values; the upper two SP FP values are passed through from *a*.

```
r0 := (float)b0
```

```
r1 := (float)b1
```

```
r2 := a2
```

```
r3 := a3
```

```
__m128 _mm_cvtpi16_ps(__m64 a)(composite)
```

Convert the four 16-bit signed integer values in *a* to four single precision FP values.

```
r0 := (float)a0
```

```
r1 := (float)a1
```

```
r2 := (float)a2
```

```
r3 := (float)a3
```

```
__m128 _mm_cvtpul6_ps(__m64 a)(composite)
```

Convert the four 16-bit unsigned integer values in **a** to four single precision FP values.

```
r0 := (float)a0
r1 := (float)a1
r2 := (float)a2
r3 := (float)a3
```

```
__m128 _mm_cvtpi8_ps(__m64 a)(composite)
```

Convert the lower four 8-bit signed integer values in **a** to four single precision FP values.

```
r0 := (float)a0
r1 := (float)a1
r2 := (float)a2
r3 := (float)a3
```

```
__m128 _mm_cvtpu8_ps(__m64 a) (composite)
```

Convert the lower four 8-bit unsigned integer values in **a** to four single precision FP values.

```
r0 := (float)a0
r1 := (float)a1
r2 := (float)a2
r3 := (float)a3
```

```
__m128 _mm_cvtpi32x2_ps(__m64 a, __m64 b) (composite)
```

Convert the two 32-bit signed integer values in **a** and the two 32-bit signed integer values in **b** to four single precision FP values.

```
r0 := (float)a0
r1 := (float)a1
r2 := (float)b0
r3 := (float)b1
```

```
__m64 _mm_cvtps_pi16(__m128 a) (composite)
```

Convert the four single precision FP values in **a** to four signed 16-bit integer values.

```
r0 := (short)a0
r1 := (short)a1
```



```
r2 := (short)a2
r3 := (short)a3
```

```
__m64 _mm_cvtps_pi8(__m128 a) (composite)
```

Convert the four single precision FP values in *a* to the lower four signed 8-bit integer values of the result.

```
r0 := (char)a0
r1 := (char)a1
r2 := (char)a2
r3 := (char)a3
```

Miscellaneous

```
__m128 _mm_shuffle_ps(__m128 a, __m128 b, int i) SHUFPS
```

Selects four specific SP FP values from *a* and *b*, based on the mask *i*. The mask must be an immediate. See [“Macro Function for Shuffle”](#) in the end of this section for a description of the shuffle semantics.

```
__m128 _mm_unpackhi_ps(__m128 a, __m128 b) UNPCKHPS
```

Selects and interleaves the upper two SP FP values from *a* and *b*.

```
r0 := a2
r1 := b2
r2 := a3
r3 := b3
```

```
__m128 _mm_unpacklo_ps(__m128 a, __m128 b) UNPCKLPS
```

Selects and interleaves the lower two SP FP values from *a* and *b*.

```
r0 := a0
r1 := b0
r2 := a1
r3 := b1
```

```
__m128 _mm_loadh_pi(__m128 a, __m64 *p) MOVHPS reg, mem
```

Sets the upper two SP FP values with 64 bits of data loaded from the address *p*; the lower two values are passed through from *a*.

```
r0 := a0
r1 := a1
r2 := *p0
r3 := *p1
```

```
void _mm_storeh_pi(__m64 *p, __m128 a)      MOVHPS mem, reg
```

Stores the upper two SP FP values of *a* to the address *p*.

```
*p0 := a2
*p1 := a3
```

```
__m128 _mm_movehl_ps (__m128 a, __m128 b) MOVHLPS
```

Moves the upper 2 SP FP values of *b* to the lower 2 SP FP values of the result.

The upper 2 SP FP values of *a* are passed through to the result.

```
r3 := a3
r2 := a2
r1 := b3
r0 := b2
```

```
__m128 _mm_movelh_ps (__m128 a, __m128 b) MOVLHPS
```

Moves the lower 2 SP FP values of *b* to the upper 2 SP FP values of the result.

The lower 2 SP FP values of *a* are passed through to the result.

```
r3 := b1
r2 := b0
r1 := a1
r0 := a0
```

```
__m128 _mm_loadl_pi(__m128 a, __m64 *p) MOVLPS reg, mem
```

Sets the lower two SP FP values with 64 bits of data loaded from the address *p*; the upper two values are passed through from *a*.

```
r0 := *p0
r1 := *p1
r2 := a2
r3 := a3
```

```
void _mm_storel_pi(__m64 *p, __m128 a)  MOVLPS mem, reg
```

Stores the lower two SP FP values of *a* to the address *p*.

```
*p0 := b0
```

```
*p1 := b1
```

```
int _mm_movemask_ps(__m128 a)
```

MOVMSKPS

Creates a 4-bit mask from the most significant bits of the four SP FP values.

```
r := sign(a3)<<3 | sign(a2)<<2 | sign(a1)<<1 | sign(a0)
```

```
unsigned int _mm_getcsr(void)
```

STMXCSR

Returns the contents of the control register.

```
void _mm_setcsr(unsigned int i)
```

LDMXCSR

Sets the control register to the value specified.

Macro Function for Shuffle

The Streaming SIMD Extensions provide a macro function to help create constants that describe shuffle operations. The macro takes four small integers (in the range of 0 to 3) and combines them into an 8-bit immediate value used by the **SHUFPS** instruction. See [Example 13-1](#).

Example 13-1 Shuffle Function Macro


```
_MM_SHUFFLE(z,y,x,w)
```

expands to the value of

```
(z<<6) | (y<<4) | (x<<2) | w
```

You can view the four integers as selectors for choosing which two words from the first input operand and which two words from the second are to be put into the result word.

Example 13-2 View of Original and Result Words With Shuffle Function Macro



Macro Functions to Read and Write the Control Registers

The following macro functions enable you to read and write bits to and from the control register. For details, see [“Set Operations”](#) later in this chapter.

Exception State Macros

```

_MM_SET_EXCEPTION_STATE(x)
_MM_GET_EXCEPTION_STATE()

```

Macro Definitions

Write to and read from the sixth-least significant control register bit, respectively.

Macro Arguments

```

_MM_EXCEPT_INVALID
_MM_EXCEPT_DIV_ZERO
_MM_EXCEPT_DENORM
_MM_EXCEPT_OVERFLOW
_MM_EXCEPT_UNDERFLOW
_MM_EXCEPT_INEXACT

```

[Example 13-3](#) tests for a divide-by-zero exception.

Example 13-3 Exception State Macros with `_MM_EXCEPT_DIV_ZERO`

```

if (_MM_GET_EXCEPTION_STATE(x) & _MM_EXCEPT_DIV_ZERO) {
    /* Exception has occurred */
}

```

Exception Mask Macros

```

_MM_SET_EXCEPTION_MASK(x)

```

Macro Arguments

```

_MM_MASK_INVALID

```

```
_MM_GET_EXCEPTION_MASK( )
```

Macro Definitions

Write to and read from the seventh through twelfth control register bits, respectively.

Note: All six exception mask bits are always affected. Bits not set explicitly are cleared.

[Example 13-4](#) masks the overflow and underflow exceptions and unmask all other exceptions.

Example 13-4 Exception Mask with `_MM_MASK_OVERFLOW` and `_MM_MASK_UNDERFLOW`

```
_MM_SET_EXCEPTION_MASK( _MM_MASK_OVERFLOW | _MM_MASK_UNDERFLOW )
```

```
_MM_MASK_DIV_ZERO
```

```
_MM_MASK_DENORM
```

```
_MM_MASK_OVERFLOW
```

```
_MM_MASK_UNDERFLOW
```

```
_MM_MASK_INEXACT
```

Rounding Mode

```
_MM_SET_ROUNDING_MODE( x )
```

```
_MM_GET_ROUNDING_MODE( )
```

Macro Definition

Write to and read from bits thirteen and fourteen of the control register.

[Example 13-5](#) tests the rounding mode for round toward zero.

Example 13-5 Rounding Mode with `_MM_ROUND_TOWARD_ZERO`

```
if ( _MM_GET_ROUNDING_MODE() == _MM_ROUND_TOWARD_ZERO ) {  
    /* Rounding mode is round toward zero */  
}
```

Macro Arguments

```
_MM_ROUND_NEAREST
```

```
_MM_ROUND_DOWN
```

```
_MM_ROUND_UP
```

```
_MM_ROUND_TOWARD_ZERO
```

Flush-to-Zero Mode

```
_MM_SET_FLUSH_ZERO_MODE( x )
```

```
_MM_GET_FLUSH_ZERO_MODE( )
```

Macro Arguments

```
_MM_FLUSH_ZERO_ON
```

```
_MM_FLUSH_ZERO_OFF
```

Macro Definition
Write to and read from bit fifteen of the control register.

Example 13-6 disables flush-to-zero mode.

Example 13-6 Flush-to-Zero Mode with `_MM_FLUSH_ZERO_OFF`

`_MM_SET_FLUSH_ZERO_MODE(_MM_FLUSHZERO_OFF)`

Macro Function for Matrix Transposition

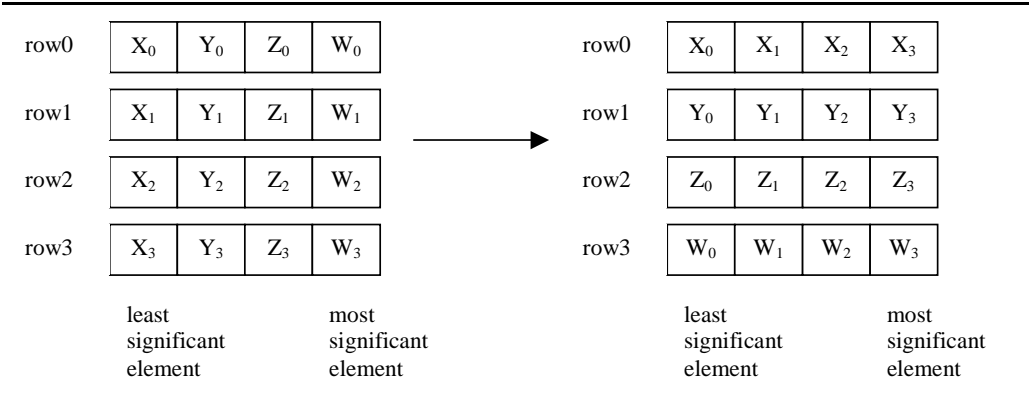
The Streaming SIMD Extensions also provide the following macro function to transpose a 4 by 4 matrix of single precision floating point values.

```
_MM_TRANSPOSE4_PS(row0, row1, row2, row3)
```

The arguments `row0`, `row1`, `row2`, and `row3` are `__m128` values whose elements form the corresponding rows of a 4 by 4 matrix. The matrix transposition is returned in arguments `row0`, `row1`, `row2`, and `row3` where `row0` now holds `column 0` of the original matrix, `row1` now holds `column 1` of the original matrix, and so on.

The transposition function of this macro is illustrated in Figure 13-2.

Figure 13-2 Matrix Transposition Using the `_MM_TRANSPOSE4_PS` Macro



Memory and Initialization

This section describes the Load, Set, and Store operations, which let you load and store data into memory. The Load and Set operations are similar in that both initialize `__m128` data. However, the Set operations take a float argument and are intended for initialization with constants, whereas the Load operations take a floating point argument and are intended to mimic the instructions for loading data from memory. The Store operation assigns the initialized data to the address.

Load Operations

`__m128 _mm_load_ss(float *p)` MOVSS

Loads an SP FP value into the low word and clears the upper three words.

```
r0 := *p
r1 := 0.0 ; r2 := 0.0 ; r3 := 0.0
```

`__m128 _mm_load1_ps(float *p)` MOVSS + shuffling

or

`__m128 _mm_load_ps1(float *p)` MOVSS + shuffling

Loads a single SP FP value, copying it into all four words.

```
r0 := *p
r1 := *p
r2 := *p
r3 := *p
```

`__m128 _mm_load_ps(float *p)` MOVAPS

Loads four SP FP values. The address must be 16-byte-aligned.

```
r0 := p[0]
r1 := p[1]
r2 := p[2]
r3 := p[3]
```

`__m128 _mm_loadu_ps(float *p)` MOVUPS

Loads four SP FP values. The address need not be 16-byte-aligned.

```
r0 := p[0]
r1 := p[1]
r2 := p[2]
r3 := p[3]
```

```
__m128 _mm_loadr_ps(float *p)          MOVAPS + shuffling
```

Loads four SP FP values in reverse order. The address must be 16-byte-aligned.

```
r0 := p[3]
r1 := p[2]
r2 := p[1]
r3 := p[0]
```

Set Operations

```
__m128 _mm_set_ss(float w)              (composite)
```

Sets the low word of an SP FP value to *w* and clears the upper three words.

```
r0 := w
r1 := r2 := r3 := 0.0
```

```
__m128 _mm_set1_ps(float w)             (composite)
```

or

```
__m128 _mm_set_ps1(float w)             (composite)
```

Sets the four SP FP values to *w*.

```
r0 := r1 := r2 := r3 := w
```

```
__m128 _mm_set_ps(float z, float y, float x, float w)
                                                    (composite)
```

Sets the four SP FP values to the four inputs.

```
r0 := w
r1 := x
r2 := y
r3 := z
```



```
__m128 _mm_setr_ps(float z, float y, float x, float w)
                                                    (composite)
```

Sets the four SP FP values to the four inputs in reverse order.

```
r0 := z
r1 := y
r2 := x
r3 := w
```

```
__m128 _mm_setzero_ps(void)
                                                    (composite)
```

Clears the four SP FP values.

```
r0 := r1 := r2 := r3 := 0.0
```

Store Operations

```
void _mm_store_ss(float *p, __m128 a)
                                                    MOVSS
```

Stores the lower SP FP value.

```
*p := a0
```

```
void _mm_storel_ps(float *p, __m128 a)
                                                    MOVSS + shuffling
```

or

```
void _mm_store_ps1(float *p, __m128 a)
                                                    MOVSS + shuffling
```

Stores the lower SP FP value across four words.

```
p[0] := a0
p[1] := a0
p[2] := a0
p[3] := a0
```

```
void _mm_store_ps(float *p, __m128 a)
                                                    MOVAPS
```

Stores four SP FP values. The address must be 16-byte-aligned.

```
p[0] := a0
p[1] := a1
p[2] := a2
p[3] := a3
```

```
void _mm_storeu_ps(float *p, __m128 a) MOVUPS
```

Stores four SP FP values. The address need not be 16-byte-aligned.

```
p[0] := a0
p[1] := a1
p[2] := a2
p[3] := a3
```

```
void _mm_storer_ps(float *p, __m128 a) MOVAPS + shuffling
```

Stores four SP FP values in reverse order. The address must be 16-byte-aligned.

```
p[0] := a3
p[1] := a2
p[2] := a1
p[3] := a0
```

```
__m128 _mm_move_ss(__m128 a, __m128 b) MOVSS
```

Sets the low word to the SP FP value of *b*. The upper 3 SP FP values are passed through from *a*.

```
r0 := b0
r1 := a1
r2 := a2
r3 := a3
```

Integer Intrinsics

The integer intrinsics are listed in [Table 13-7](#), followed by a description of each intrinsic with the most recent mnemonic naming convention. The alternate name is provided in case you have used these intrinsics before.

Table 13-7 Integer Intrinsics

Intrinsic Name	Alternate Name	Corresponding Instruction
<code>_mm_extract_pi16</code>	<code>_m_pextrw</code>	<code>PEXTRW</code>
<code>_mm_insert_pi16</code>	<code>_m_pinsrw</code>	<code>PINSRW</code>
<code>_mm_max_pi16</code>	<code>_m_pmaxsw</code>	<code>PMAXSW</code>
<code>_mm_max_pu8</code>	<code>_m_pmaxub</code>	<code>PMAXUB</code>

Table 13-7 Integer Intrinsics (Continued)

Intrinsic Name	Alternate Name	Corresponding Instruction
<code>_mm_min_pi16</code>	<code>_m_pminsw</code>	PMINSW
<code>_mm_min_pu8</code>	<code>_m_pminub</code>	PMINUB
<code>_mm_movemask_pi8</code>	<code>_m_pmovmskb</code>	PMOVMASKB
<code>_mm_mulhi_pu16</code>	<code>_m_pmulhuw</code>	PMULHUW
<code>_mm_shuffle_pi16</code>	<code>_m_pshufw</code>	PSHUFW
<code>_mm_maskmove_si64</code>	<code>_m_maskmovq</code>	MASKMOVQ
<code>_mm_avg_pu8</code>	<code>_m_pavgb</code>	PAVGB
<code>_mm_avg_pu16</code>	<code>_m_pavgw</code>	PAVGW
<code>_mm_sad_pu8</code>	<code>_m_psadbw</code>	PSADBW

For this section you need to ensure to empty the multimedia state for the `mmx` register. See [“The EMMS Instruction: Why You Need it and When to Use it”](#).

```
int _mm_extract_pi16(__m64 a, int n) PEXTRW
```

Extracts one of the four words of `a`. The selector `n` must be an immediate.

```
r := (n==0) ? a0 : ( (n==1) ? a1 : ( (n==2) ? a2 : a3 ) )
```

```
__m64 _mm_insert_pi16(__m64 a, int d, int n) PINSRW
```

Inserts word `d` into one of four words of `a`. The selector `n` must be an immediate.

```
r0 := (n==0) ? d : a0;
```

```
r1 := (n==1) ? d : a1;
```

```
r2 := (n==2) ? d : a2;
```

```
r3 := (n==3) ? d : a3;
```

```
__m64 _mm_max_pi16(__m64 a, __m64 b) PMAXSW
```

Computes the element-wise maximum of the words in `a` and `b`.

```
r0 := min(a0, b0)
```

```
r1 := min(a1, b1)
```

```
r2 := min(a2, b2)
```

```
r3 := min(a3, b3)
```

```
__m64 _mm_max_pu8(__m64 a, __m64 b) PMAXUB
```

Computes the element-wise maximum of the unsigned bytes in *a* and *b*.

```
r0 := min(a0, b0)
r1 := min(a1, b1)
...
r7 := min(a7, b7)
```

```
__m64 _mm_min_pi16(__m64 a, __m64 b) PMINSW
```

Computes the element-wise minimum of the words in *a* and *b*.

```
r0 := min(a0, b0)
r1 := min(a1, b1)
r2 := min(a2, b2)
r3 := min(a3, b3)
```

```
__m64 _mm_min_pu8(__m64 a, __m64 b) PMINUB
```

Computes the element-wise minimum of the unsigned bytes in *a* and *b*.

```
r0 := min(a0, b0)
r1 := min(a1, b1)
...
r7 := min(a7, b7)
```

```
int _mm_movemask_pi8(__m64 a) PMOVBMSKB
```

Creates an 8-bit mask from the most significant bits of the bytes in *a*.

```
r := sign(a7)<<7 | sign(a6)<<6 | ... | sign(a0)
```

```
__m64 _mm_mulhi_pu16(__m64 a, __m64 b) PMULHUW
```

Multiplies the unsigned words in *a* and *b*, returning the upper 16 bits of the 32-bit intermediate results.

```
r0 := hiword(a0 * b0)
r1 := hiword(a1 * b1)
r2 := hiword(a2 * b2)
r3 := hiword(a3 * b3)
```

```
__m64 _mm_shuffle_pi16(__m64 a, int n) PSHUFW
```

Returns a combination of the four words of *a*. The selector *n* must be an immediate.

```
r0 := word (n&0x3) of a
r1 := word ((n>>2)&0x3) of a
r2 := word ((n>>4)&0x3) of a
r3 := word ((n>>6)&0x3) of a
```

```
void _mm_maskmove_si64(__m64 d, __m64 n, char *p)
MASKMOVQ
```

Conditionally store byte elements of *d* to address *p*. The high bit of each byte in the selector *n* determines whether the corresponding byte in *d* will be stored.

```
if (sign(n0)) p[0] := d0
if (sign(n1)) p[1] := d1
...
if (sign(n7)) p[7] := d7
```

```
__m64 _mm_avg_pu8(__m64 a, __m64 b) PAVGB
```

Computes the (rounded) averages of the unsigned bytes in *a* and *b*.

```
t = (unsigned short)a0 + (unsigned short)b0
r0 = (t >> 1) | (t & 0x01)
...
t = (unsigned short)a7 + (unsigned short)b7
r7 = (unsigned char)((t >> 1) | (t & 0x01))
```

```
__m64 _mm_avg_pu16(__m64 a, __m64 b) PAVGW
```

Computes the (rounded) averages of the unsigned words in *a* and *b*.

```
t = (unsigned int)a0 + (unsigned int)b0
r0 = (t >> 1) | (t & 0x01)
...
t = (unsigned word)a7 + (unsigned word)b7
r7 = (unsigned short)((t >> 1) | (t & 0x01))
```

```
__m64 _mm_sad_pu8(__m64 a, __m64 b) PSADBW
```

Computes the sum of the absolute differences of the unsigned bytes in **a** and **b**, returning the value in the lower word. The upper three words are cleared.

```
r0 = abs(a0-b0) + ... + abs(a7-b7)
r1 = r2 = r3 = 0
```

Cacheability Support

The following intrinsics provide ways to make efficient use of the cache.

```
void _mm_prefetch(char *p, int i) PREFETCH
```

Loads one cache line of data from address **p** to a location “closer” to the processor. The value **i** specifies the type of prefetch operation: the constants `_MM_HINT_T0`, `_MM_HINT_T1`, `_MM_HINT_T2`, and `_MM_HINT_NTA` should be used, corresponding to the type of prefetch instruction.

```
void _mm_stream_pi(__m64 *p, __m64 a) MOVNTQ
```

Stores the data in **a** to the address **p** without polluting the caches. This intrinsic requires you to empty the multimedia state for the **mmx** register. See [“The EMMS Instruction: Why You Need it and When to Use it”](#).

```
void _mm_stream_ps(float *p, __m128 a) MOVNTPS
```

Stores the data in **a** to the address **p** without polluting the caches. The address must be 16-byte-aligned.

```
void _mm_sfence(void) SFENCE
```

Guarantees that every preceding store is globally visible before any subsequent store.

Data Alignment

To improve performance, you should align data to 16 bytes in memory operations when you are using the Streaming SIMD Extensions.

Specifically, you must align `__m128` objects as addresses passed to the `_mm_load` and `_mm_store` intrinsics. If you want to declare arrays of floats and treat them as `__m128` objects by casting, you need to ensure that the float arrays are properly aligned.

Alignment Support

Use `__declspec(align)` to direct the compiler to align data more strictly than it otherwise does. For example, a data object of type `int` is allocated at a byte address which is a multiple of 4 by default (the size of an `int`). However, by using `__declspec(align)`, you can direct the compiler to instead use an address which is a multiple of 8, 16, or 32 with the following restrictions:

- 32-byte addresses must be statically allocated
- 16-byte addresses can be locally or statically allocated

You can use this data alignment support as an advantage in optimizing cache line usage. By clustering small objects that are commonly used together into a `struct`, and forcing the `struct` to be allocated at the beginning of a cache line, you can effectively guarantee that each object is loaded into the cache as soon as any one is accessed, resulting in a significant performance benefit.

The syntax of this extended-attribute is as follows:

`align(n)`

where `n` is an integral power of 2, less than or equal to 32. The value specified is the requested alignment.



NOTE. *If a value is specified that is less than the alignment of the affected data type, it has no effect. In other words, data is aligned to the maximum of its own alignment or the alignment specified with `__declspec(align)`.*

You can request alignments for individual variables, whether of static or automatic storage duration. (Global and static variables have static storage duration; local variables have automatic storage duration by default.)

However, you cannot adjust the alignment of a parameter, nor a field of a `struct` or `class`. However, you can increase the alignment of a `struct` (or `union` or `class`), in which case every object of that type is affected.

As an example, suppose that a function uses local variables `i` and `j` as subscripts into a 2-dimensional array. They might be declared as follows:

```
int i, j;
```

These variables are commonly used together. But they can fall in different cache lines, which could be detrimental to performance. You can instead declare them as follows:

```
__declspec(align(8)) struct { int i, j; } sub;
```

The compiler now ensures that they are allocated in the same cache line. In C++, you can omit the `struct` variable name (written as `sub` in the above example). In C, however, it is required, and you must write references to `i` and `j` as `sub.i` and `sub.j`.

If you use many functions with such subscript pairs, it is more convenient to declare and use a struct type for them, as in the following example:

```
typedef struct __declspec(align(8)) { int i, j; } Sub;
```

By placing the `__declspec(align)` after the keyword `struct`, you are requesting the appropriate alignment for all objects of that type. However, that allocation of parameters is unaffected by `__declspec(align)`. If necessary, you can assign the value of a parameter to a local variable with the appropriate alignment.)

You can also force alignment of global variables, such as arrays:

```
__declspec(align(16)) float array[1000];
```

Dynamic Stack Frame Alignment

By default, the compiler targets functions that use 8- or 16-byte objects for dynamic alignment. It is up to the discretion of the compiler to decide whether dynamic stack alignment is needed on a function-per-function level (for stability of performance). This ensures that references to `double` and `__m64` to `double`, `__m64`, and `__m128` local variables align to addresses evenly divisible by 16. Parameters alignment is not affected.

Table 13-8 shows how `-Qsalign` controls the stack frame alignment.

Table 13-8 Stack Frame Alignment Options

Option	Meaning
<code>-Qsalign8</code>	Align stack for frames or functions with variables of 8 or 16 bytes where it is deemed appropriate by the compiler.
<code>-Qsalign16</code>	Align stack for frames or functions with 16-byte variables where it is deemed appropriate by the compiler.
<code>-Qsalign-</code>	Disable stack alignment for all functions.



If you are using aligned frames, you should not modify the `EBX` register in inlined assembly blocks because `EBX` is used to keep track of the argument block. You can modify `EBX` only if you save and restore `EBX` each time you use it. The Intel C/C++ Compiler uses the `EBX` register to control alignment of variables of these data types so using `EBX`, without preserving it, will cause unexpected program execution..

For additional information on the use of `EBX` in inline assembly code and other related issues, see the Application Note AP-833 *Data Alignment and Programming Issues with the Intel C/C++ Compiler* (Order #243872-001), and AP-589 *Software Conventions for the Streaming SIMD Extensions* (Order # 243873-001).

Assembly Language Support

The Intel C/C++ Compiler supports inlining of assembly blocks and generation of assembly files.

Inline Assembly

The compiler supports use of all the MMX and Streaming SIMD Extensions in inline assembly (`__asm`) blocks. The compiler also accepts the new syntax `MMWORD PTR` and `XMMWORD PTR` to refer to 64- and 128-bit data.

Generation of Assembly files

Use the `-S` compiler switch to produce an assembly listing. This is useful for hand-tuning compiler-generated code. The assembly files can be directly compiled using MASM version 6.12 with `iaxmm.inc` MASM include files. Generated `.asm` files contain an include directive for them when the MMX technology intrinsics or Streaming SIMD Extensions are used. For details on using inline assembly, see [“Producing an Assembly Code Listing \(-S\)” in Chapter 6](#).

Compiler Vectorization Support and Guidelines

14

This chapter provides guidelines, option descriptions, and examples for compiler vectorization with respect to the Intel C/C++ Compiler. The following list provides a summary of the chapter contents.

- A Quick reference of vectorization functionality and features
- Descriptions of compiler switches to control vectorization
- Descriptions of the C/C++ language features to control vectorization
- Discussion and general guidelines on vectorization levels:
 - automatic vectorization
 - vectorization with user intervention
 - no vectorization whatsoever
- Examples demonstrating typical vectorization issues and resolutions

Vectorizer Quick Reference

The Intel C/C++ Compiler vectorization can be summarized by the command line switches in [Table 14-1](#):

Table 14-1 **Vectorization Command-Line Switches**

Option	Description	Default	Reference
<code>-Qvec</code>	Enables the vectorizer.	OFF	page 14-2
<code>-Qvec_alignment</code>	Controls the default alignment of vectorizable data.	OFF	page 14-2
<code>-Qvec_verbose</code>	Controls the vectorizer's diagnostic levels.	<code>n=1</code>	page 14-3

continued

Table 14-1 Vectorization Command-Line Switches (Continued)

Option	Description	Default	Reference
<code>-Qrestrict</code>	Enables pointer disambiguation with the <code>restrict</code> qualifier.	OFF	page 14-3
<code>-Qkscalar</code>	Performs all 32-bit floating point arithmetic using the Streaming SIMD Extensions instead of the default x87 instructions.	OFF	page 14-3
<code>-Qvec_emms[-]</code>	Controls the automation of <code>EMMS</code> instruction insertions to empty the <code>mmx</code> instruction registers.	ON	page 14-3
<code>-Qvec_no_arg_alias[-]</code>	Assumes on entry that procedure arguments are not aliased.	OFF	page 14-4
<code>-Qvec_no_alias[-]</code>	Assumes that no aliasing can occur between objects with different names.	OFF	page 14-4

Command-Line Switch Support

The following vectorizer switch options are available from the compiler command line. As with many advanced features of compilers, you must be sure to properly understand the functionality of the switches in order to use them effectively and avoid unwanted program behavior.

`-Qvec` Enables the vectorizer. All other vectorization switches require this switch to on. The default setting is off .

`-Qvec_alignment[n]` Controls the default alignment of data that can be vectorized. By default, the compiler assumes a natural alignment; for example, a `float` array aligned to 4 bytes generates *unaligned* memory operations. However, if the compiler proves proper alignment, then it uses aligned instructions to improve performance. To override the compiler's default alignment, use the sub options as follows:

`n=0`: default alignment
`n=1`: assumes unaligned
`n=2`: assumes aligned

This feature is commonly used in programs that do not use `__declspec(align)` to pass data with pointers. See [“Data Alignment” in Chapter 13](#).



CAUTION. *If you change the default alignment, you must be absolutely sure . Otherwise, the compiler will generate incorrect code.*

`-Qvec_emms[-]`

Controls whether the compiler inserts an **EMMS** instruction after each integer loop vectorized. See [“The EMMS Instruction: Why You Need it and When to Use it” in Chapter 13](#).



CAUTION. *If you cannot tell whether your program uses MMX™ technology, only use `-Qvec_emms[-]` when you know it is safe to do so.*

`-Qrestrict`

The **restrict** keyword is a qualifier that allows the disambiguator flexibility in alias assumptions, which enables more vectorization and other optimizations. Because the **restrict** keyword is a language extension, you must use `-Qrestrict` to enable it.

`-Qkscalar`

Specifies 32-bit floating point arithmetic to operate on 32-bit scalar operations with Streaming SIMD Extensions.

`-Qvec_verbose[n]`

Controls the vectorizer’s level of diagnostic messages:

- n=0:** no diagnostic information is displayed
- n=1:** display diagnostics indicating loops successfully vectorized (default)
- n=2:** same as **n=1**, plus diagnostics indicating loops not successfully vectorized

`n=3:` same as `n=2`, plus additional information about any proven or assumed dependences

`-Qvec_no_arg_alias[-]`

Assumes that procedure arguments are not aliased on entry, neither by direct nor indirect referencing. The loop in [Example 14-1](#) will not vectorize unless you use the `-Qvec_no_arg_alias` switch.

Example 14-1 Using `-Qvec_no_arg_alias` to Vectorize

```
void f(float *out, float *in){
    int i;
    for (i = 0; i < 100; i++)
        out[i] = in[i];
}
```



CAUTION. *If you incorrectly assert that procedure arguments are not aliased, you will generate improper program behavior. This caution applies to `-Qvec_no_arg_alias` and `-Qvec_no_alias`.*

`-Qvec_no_alias[-]`

Assumes that no aliasing occurs between objects with different names, neither by direct nor indirect referencing. The loop in [Example 14-2](#) does not vectorize unless you use the `-Qvec_no_alias`.

Example 14-2 Using `-Qvec_no_alias` to Vectorize

```

void f(float *p){
    int i;
    float *out=p;
    float *in=p;

    for (i = 0; i < 100; i++)
        out[i] = in[i];
}

```

Language Support and Pragmas

This section addresses language features that help vectorize to the target hardware limitations and stylistic aspects. The `declspec(align(n))` declaration enables you to overcome hardware alignment constraints. The `restrict` qualifier and the pragmas address the stylistic issues due to lexical scope, data interdependence, and ambiguity resolution.

Table 14-2 Language Support

Option	Description	Reference
<code>declspec(align(n))</code>	Directs the compiler to align the variable <code>var-name</code> to an <code>n</code> -byte boundary.	page 14-6
<code>restrict</code>	Permits the disambiguator flexibility in alias assumptions, which enables more vectorization.	page 14-3
<code>#pragma ivdep</code>	Instructs the compiler to ignore assumed vector dependencies.	page 14-7
<code>#pragma vector {aligned unaligned}</code>	Specifies how to vectorize the loop.	page 14-8
<code>#pragma novector</code>	Specifies that the loop should never be vectorized	page 14-9

Alignment with declspec

Syntax: `__declspec(align(n)) <type> var-name`

Definition: Directs the compiler to align the variable *var-name* to an *n*-byte boundary, where *n* must be between 2 and 32. See [“Data Alignment” in Chapter 13](#).

[Example 14-3](#) directs the compiler to align the `float` array to 16 bytes.

Example 14-3 Using `declspec(align(n))` for Alignment

```
__declspec(align(16)) float x[1024];
```

Qualifying with the restrict Keyword

Syntax: `<type> *restrict p`

Definition: The `restrict` keyword is a qualifier. If *p* is a restricted pointer to an object *x*, then all references to *x* within the lexical scope of *p* must involve an expression with *p*. This allows the disambiguator flexibility in alias assumptions, which enables more vectorization and other optimizations. Typically, this feature is used to indicate that arguments to the function are not aliases.

Because the `restrict` keyword is a language extension, you must use `-Qrestrict` to enable it. Without the restricted pointers in the following, the compiler must assume that `a[i]` is aliased to `b[i-1]`, which prohibits vectorization.

Example 14-4 Using the restrict Keyword as a Qualifier

```
void foo (float *restrict a, float *b) {  
    for (i=0; i<N; i++)  
        a[i] = b[i] * 2.0f;  
}
```

Pragma Scope

These pragmas control the vectorization of only the subsequent loop in the program, but the program does not apply them to any nested loops. Each nested loop needs its own pragma preceding it in order for the pragma to be applied. You must place a pragma only before the loop control statement. However, assignments and function calls can appear between the pragma and the loop it affects.

#pragma ivdep

Syntax: `#pragma ivdep`

Definition: This pragma instructs the compiler to ignore assumed vector dependences. To ensure correct code, the compiler treats an assumed dependence as a proven dependence, which prevents vectorization. This pragma overrides that decision. Only use this when you know that the assumed loop dependences are safe to ignore.

The loop in [Example 14-5](#) will not vectorize without the `ivdep` pragma, since `k` is not known (and it could be `k<0`).

Example 14-5 Loop Using #pragma ivdep

```
#pragma ivdep
for (i=0; i<m; i++) {
    a[i] = a[i+k] * c;
    ...
}
```

#pragma vector

Syntax: `#pragma vector {aligned |unaligned}`

Definition: The `vector` loop pragma means the loop should be vectorized, if it is legal to do so, ignoring normal heuristic decisions about profitability. When the `aligned` (or `unaligned`) qualifier is used with this pragma, the loop should be vectorized using `aligned` (or `unaligned`) operations. Specify one and only one of `aligned` or `unaligned`. Using pragma to control alignment in a loop overrides `-Qvec_alignment[n]`.

The loop in [Example 14-6](#) uses the vector pragma to indicate that, despite the stride-2 references, the loop should be vectorized.

Example 14-6 Loop Using Stride-2 #pragma vector

```
#pragma vector
for (i=0; i<m; i+=2) {
    a[i] = a[i] * c;
}
```

This loop in [Example 14-7](#) uses the `aligned` qualifier to request that the loop be vectorized with aligned instructions, as the arrays are declared in such a way that the compiler could not normally prove this would be safe to do so.

Example 14-7 Loop Using #pragma vector aligned

```
void f(float *a) {  
    #pragma vector aligned  
    for (i=0; i<m; i+=2) {  
        a[i] = a[i] * c;  
    }  
}
```

#pragma novector

Syntax: `#pragma novector`

Definition: The `novector` loop pragma specifies that the loop should never be vectorized, even if it is legal and profitable to do so.

In example [Example 14-8](#), the `novector` pragma indicates that the programmer knows the tripcount is too low to make vectorization be worthwhile. This is indicated by the value of (`ub - lb`) being small.

Example 14-8 Low Trip Count Loop Using #pragma novector

```
void f(int lb, int ub) {  
    ...  
    #pragma novector  
    for (j=lb; j<ub; j++) {  
        a[j] = a[j] + b[i];  
    }  
    ...  
}
```

Loop Structure Coding Background

The goal of vectorizing compilers is to parallelize source code automatically. However, the realization of this goal has been difficult to achieve even for well-known vectorizing compilers. The reason for the difficulty in achieving automatic parallelization is due to two major factors:

1. *Style*—The style in which you write source code can inhibit optimization. For example, a common problem with global pointers is that they often prevent the compiler from being able to prove two memory references are distinct locations. Consequently, this prevents certain reordering transformations.
2. *Hardware Restrictions*—The compiler is limited by restrictions imposed by the underlying hardware. In the case of Streaming SIMD Extensions, the vector memory operations are limited to stride-1 accesses with a preference to 16-byte aligned memory references. This means that if the compiler abstractly recognizes a loop as vectorizable, it still might not vectorize it to a distinct target architecture.

Many stylistic issues that prevent the automatic parallelization by vectorization compilers are found in loop structures. The ambiguity arises from the complexity of the keywords, operators, data references, and memory operations within the loop bodies.

However, by understanding these limitations and by knowing how to interpret diagnostic messages, you can modify your program to overcome the historical difficulty and enable optimal vectorizations. The following sections summarize the capabilities and restrictions of the vectorizer with respect to loop structures.

Key Programming Guidelines

Review these guidelines, restrictions, and examples, and check them against your code to eliminate ambiguities that prevent the compiler from achieving optimal vectorization.

- Guidelines for loop bodies:
 - Use straight-line code (a single basic block)
 - Use vector data only; that is, arrays, and invariant expressions on the right hand side of assignments. Array references can appear on the left hand side of assignments.

- Use only assignment statements
- *Avoid* the following in loop bodies:
 - Function calls
 - Unvectorizable operations
 - Mixing vectorizable types in the same loop
 - Multiple types of SIMD operations
 - Use relational operators or the `?:` operator
 - Data-dependent loop exit conditions

Loop Constructs

Loops can be formed with the usual `for` and `while-do`, or repeat-until constructs or by using a `goto` or a label. However, the loops must have a single entry and a single exit to be vectorized.

Example 14-9 Loop Construct Usage

Correct Usage	Incorrect Usage
<pre>for (i=0; i<n; i++) { ... }</pre>	<pre>while (i<n) { ... if (cond) break; /* 2nd exit */ ... }</pre>

Loop Body Control Flow

Loop bodies must be straight-line code, that is, no branches.

Example 14-10 Loop Body Control Flow

Correct Usage:

```
do {  
    /* body is a single basic block */  
    a[i] = b[i] * x  
    b[i] = c[i] + sqrt(d[i]);  
    --count;  
} while (count != 0);
```

Incorrect Usage:

```
while (i<n) {  
    /* if-branch inside body of loop */  
    a[i] = b[i] * c[i];  
    if (a[i] < 0.0) {  
        a[i] = 0.0;  
    }  
}
```

Loop Exit Conditions

Loop exit conditions determine the number of iterations that a loop executes. For example, fixed indexes in **for** loops determine the iterations. The loop iterations must be countable; that is, the number of iterations must be expressed as one of the following:

- a constant
- a linear function of an integer variable
- a loop invariant terms

Loops whose exit depends on computation are not countable.

Example 14-11 Loop Usage Comparisons**Correct Usage for Countable Loop:**

```
count = N; /* exit condition specified by "N - lb + 1" */
...
do {
  a[i] = b[i] * x
  b[i] = c[i] + sqrt(d[i]);
  --count;
} while (count != lb); /* lb is not defined within loop */
```

Correct Usage for Countable Loop:

```
/* exit condition is "(n-m+2)/2" */
i = 0;
for (l=m; l<n; l+=2) {
  a[l] = b[l] * x
  b[l] = c[l] + sqrt(d[l]);
  ++i;
}
```

Incorrect Usage for Non-Countable Loop:

```
i = 0;
/* iterations dependent on a[i] */
while (a[i] > 0.0) {
  a[i] = b[i] * c[i];
  ++i;
}
```

Stripmining and Cleanup

The compiler will automatically stripmine your loop and generate a cleanup loop. This means you do not need to unroll your loops.

Example 14-12 Strip Mining and Cleanup Loops

```
i = 0;
while (i < n) {
    a[i] = b[i] + c[i]; /* Code Clean-up Loop. */
    ++i;
}

/* The vectorizer will generate the following two loops. */
i = 0;
while (i < (n - n%4)) {
    a[i] = b[i] + c[i]; /* Vector Loop. */
    ++i;
}
while (i < n) {
    a[i] = b[i] + c[i]; /* Scalar clean-up loop. */
    ++i;
}
```

Statements in the Loop Body

The vectorizable operations are different for floating point and integer data.

Floating-point array operations

The statements within the loop body may be `float` operations (typically on arrays). Arithmetic operations are limited to addition, subtraction, multiplication, division, and negation. Note that conversion to/from floats is not allowed. Operation on `double` types is not allowed.

Integer Array Operations

The statements within the loop body may be `char`, `unsigned char`, `short`, `unsigned short`, `int`, and `unsigned int` operations (again, typically for arrays). Arithmetic operations are limited to addition, subtraction, bitwise `AND`, `OR`, and `XOR` operators.

Other Integer Operations

You cannot mix data types except for integer induction variables.

Other Datatypes

No statements other than the preceding floating point and integer operations are allowed. In particular, note that the special `__m64` and `__m128` datatypes are not vectorizable.

No Function Calls

The loop body may not contain any function calls. The current version does not support vectorization of the math intrinsics `sqrt` and `fabs`. Use of the Streaming SIMD Extensions intrinsics (`_mm_add_ps`) are not allowed.

Vectorizable Data References

For any data reference, either as an array element or pointer reference, take care to ensure that there are no potential dependence or alias constraints preventing vectorization; intuitively, an expression in one iteration must not depend on the value computed in a previous iteration and pointer variables must provably point to distinct locations. Use of the `ivdep` pragma and the `restrict` keyword can be used to tell the compiler to ignore assumed dependences. See also the examples in the [“Data alignment”](#) section.

Arrays

Vectorizable data in a loop may be expressed as uses of array elements, provided that the array references are stride-1 (4 bytes) or loop invariant. Stride-1 references are not vectorized by default; the `vector` pragma can be used to override this.

Pointers

Vectorizable data can also be expressed using pointers, subject to the same constraints as uses of array elements: the references that are stride-1 or loop invariant. Non-unit stride references for integers are not supported.

Invariants

Vectorizable data can also include loop invariant references on the right hand inside an expression, either as variables or numeric constants. The loop in [Example 14-13](#) will vectorize:

Example 14-13 Vectorizable Loop Invariant Reference

```
for (i=0; i<n; i++) {  
    a[i] = b[i] * 3.14f + c[j];  
}
```

Data alignment

If vectorizable `float` data is provably aligned, the compiler will generate aligned instructions. This is the case for locally declared data and data declared using the alignment `declspec`. Where data alignment is not known, unaligned references will be used unless a pragma or command-line switch is used to override this as described in [“Alignment with declspec”](#).

Common Errors in Making Code Vectorization-Compatible

To make your code vectorizable, you will often need to make some changes to your loops. However, you should make *only* the changes needed to enable vectorization and no others. In particular, you should avoid these common changes:

- Do not unroll your loops, the compiler does this automatically.
- Do not decompose one loop with several statements in the body into several single-statement loops.

- Do not manually insert calls to `EMMS`; for example, via the `_m_empty` intrinsic, after the loops to be vectorized, the compiler does this by default (use `-Qvec_emms` to override the default). See `“-Qvec emms[-]”` section earlier in this chapter.

Some Examples

This section contains a few simple examples of some common issues in vector programming.

Argument Aliasing: A Vector Copy

The loop in [Example 14-14](#), a vector copy operation, does not vectorize because the compiler cannot prove `dest[i]` and `src[i]` are distinct.

Example 14-14 Unvectorizable Copy Due to Unproven Distinction

```
void vec_copy(float *dest, float *src, int len){
    int i;
    for (i=0; i<len; i++)
        dest[i] = src[i];
}
```

The `restrict` keyword indicates that the pointers refer to distinct objects. Therefore, the compiler allows vectorization.

Example 14-15 Using restrict to Prove Vectorizable Distinction

```
void vec_copy(float restrict *dest, float restrict *src, int len){
    int i;
    for (i=0; i<len; i++)
        dest[i] = src[i];
}
```

Data Alignment: Two Examples

Data Alignment Examples

The next function in [Example 14-16](#) has a loop that will vectorize but only with unaligned memory instructions because it uses global variables for which the compiler can neither prove nor force 16-byte alignment.

Example 14-16 Unaligned Loop Due to Global Variables

```
float z[N], a[N], y[N], x;
void f(void)
{
    for (i=0; i<N; i++) {
        a[i] = a[i] * x + y[i];
    }
}
```

Use the `declspec` alignment to improve the loop as in [Example 14-17](#).

Example 14-17 Loop Alignment with `declspec`

```
__declspec(align(16)) float z[N], a[N], y[N], x;
void f(void)
{
    for (i=0; i<N; i++) {
        a[i] = a[i] * x + y[i];
    }
}
```

[Example 14-18](#) has a loop that vectorizes but only with unaligned memory instructions. The compiler can align the local arrays, but because `lb` is not known at compile-time, the correct alignment cannot be determined.

Example 14-18 Loop Unaligned Due to Unknown Variable Value at Compile Time

```
void f(int lb){
    float z2[N], a2[N], y2[N], x2;
    ...
    for (i=lb; i<N; i++) {
        a2[i] = a2[i] * x2 + y2[i];
    }
    ...
}
```

If you know that `lb` is a multiple of 4, you can align the loop with `#pragma vector aligned` as shown in [Example 14-19](#).

Example 14-19 Alignment Due to Assertion of Variable as Multiple of 4

```
void f(int lb)
{
    float z2[N], a2[N], y2[N], x2;
    ...
    assert(lb%4==0);
    #pragma vector aligned
    for (i=lb; i<N; i++) {
        a2[i] = a2[i] * x2 + y2[i];
    }
    ...
}
```

The use of the assertion checks that the constraint `1b` is a multiple of 4.

Data Dependence

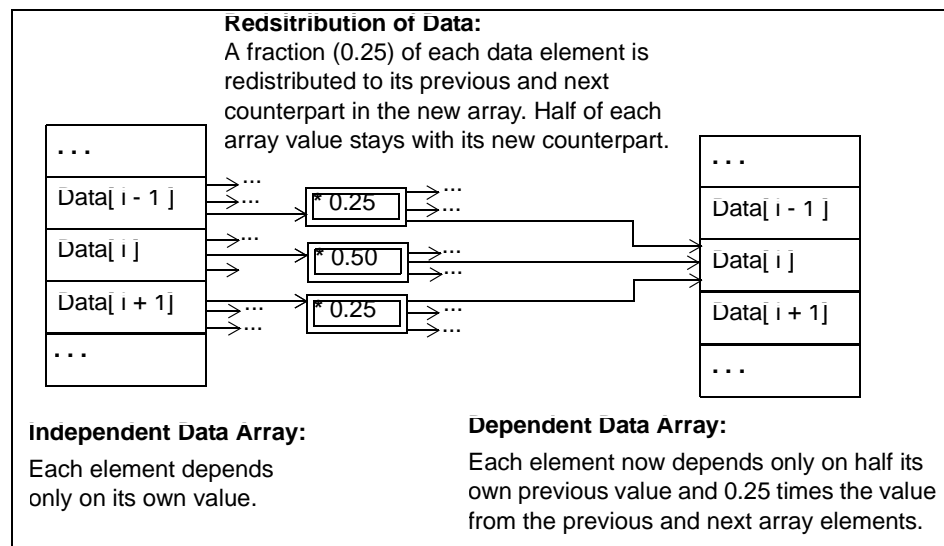
[Example 14-20](#) shows some code that exhibits data dependence. Each element of an array is changed to be function of itself and its two neighbors.

Example 14-20 Data Dependent Loop

```
float data[N];
int i;
for (i=1; i<N-1; i++) {
    data[i] = data[i-1]*0.25 + data[i]*0.5 + data[i+1]*0.25
}
```

[Figure 14-1](#) shows how [Example 14-20](#) exhibits data dependence.

Figure 14-1 A Redistribution for Data Dependence



This loop from [Example 14-20](#) is not vectorizable because the write to the current element `data[i]` is dependent on the use of the preceding element `data[i-1]`, which has already been written to and changed in the previous iteration.

To see this, look at the access patterns of the array for the first two iterations as shown in [Example 14-21](#)

Example 14-21 Data Dependence Vectorization Patterns

Unvectorizable Access Pattern

```
i=1:READ data[0]
    READ data[1]
    READ data[2]
    WRITE data[1]

i=2:READ data[1]
    READ data[2]
    READ data[3]
    WRITE data[2]
```

Vectorizable Access Pattern

```
i=1,i=2:READ data[0]
        READ data[1]
        READ data[2]
        READ data[1]
        READ data[2]
        READ data[3]
        READ data[2]
        WRITE data[1]
        WRITE data[2]
```

In the normal sequential version of this loop shown in the Unvectorizable Access Pattern, the value of `data[1]` read from during the second iteration was written to in the first iteration.

For vectorization, the iterations must be done in parallel as in the Vectorizable Access Pattern. In this version, the value of `data[1]` read is the original value, not the updated version. You might be tempted to rewrite the loop in a vectorizable style by using a temporary array followed by a copy operation as shown in [Figure 14-22](#).

Example 14-22 Data Independent but Still Unvectorizable Loop

```
float data[N], newdata[N];
int i;
for (i=1; i<N-1; i++) {
    newdata[i] = data[i-1] * 0.25 + data[i]*0.5 + data[i+1]*0.25
}
for (i=1; i<N-1; i++) {
    data[i] = newdata[i];
}
```

While both of these loops will vectorize, this computation *will produce a different result* than the original computation because you have changed the semantics of the original loop.

Loop Interchange and Subscripts: Matrix Multiply

Matrix multiplication is commonly written as shown in [Example 14-23](#).

Example 14-23 Typical Matrix Multiplication

```
for (i=0; i<N; i++) {
    for (j=0; j<n; j++) {
        for (k=0; k<n; k++) {
            c[i][j] = c[i][j] + a[i][k] * b[k][j];
        }
    }
}
```


The use of `b[k][j]`, is not a stride-1 reference and therefore will not normally be vectorizable. The use of `c[i][j]` on the left-hand side also prevents vectorization; invariant references on the left-hand side are not allowed. If the loops are interchanged, however, all the references will become stride-1 as in [Example 14-24](#).

Example 14-24 Matrix Multiplication with Stride-1

```
for (i=0; i<N; i++) {  
    for (k=0; k<n; k++) {  
        for (j=0; j<n; j++) {  
            c[i][j] = c[i][j] + a[i][k] * b[k][j];  
        }  
    }  
}
```

For Additional Information

The following sources might be useful in helping you understand basic vectorization terminology and technology:

- *High Performance Computing* (2nd edition), Kevin Dowd (O'Reilly and Associates, 1998), ISBN 156592312X
- *Intel Architecture Optimization Manual*, Intel Corporation, order number, 730795.

Compiler Limits



[Table A-1](#) shows the size or number of each item that the compiler can process. All capacities shown in the table are tested values; the actual number can be greater than the number shown.

Table A-1 **Compiler Limits**

Item	Tested Values
Control structure nesting (block nesting)	512
Conditional compilation nesting	512
Declarator modifiers	512
Parenthesis nesting levels	512
Significant characters, internal identifier	2048
External identifier name length	64K
Number of external identifiers/file	128K
Number of identifiers in a single block	2048
Number of macros simultaneously defined	128K
Number of parameters to a function call	512
Number of parameters per macro	512
Number of characters in a string	128K
Bytes in an object	512K
Include file nesting depth	512
Case labels in a switch	32K
Members in one structure or union	32K
Enumeration constants in one enumeration	8192
Levels of structure nesting	320

Experimental Performance Tuning

B

This appendix provides information on how you can adjust the compiler's optimization for a particular application by experimenting with interprocedural optimizations.



NOTE. *The suboptions described in this chapter are provided for the benefit of developers who are willing to experiment with advanced optimizations. Please be aware that experimenting with these advanced optimization can yield unexpected results. Therefore, this guide describes only the basic functionality of these advanced optimizations.*

Keywords for Optimization (-Qoption,c,optlist)

Enter the `-Qoption` option with the applicable keywords to select particular in-line expansions and loop optimizations. The option must be entered with a `-Qip` specification, as follows:

`-Qip [-Qoption,c,optlist]`

`c` is the specification that indicates the C compiler tool.

`optlist` is any applicable optimization for the specified level.

For example, the following disables the passing of arguments in registers:

```
prompt> icl -Qip -Qoption,c,ip_args_in_regs=0 a.cpp
```

The next section defines keywords to use with `-Qip` and `-Qoption`.

Interprocedural Optimization

If you specify `-Qip` without the `-Qoption` qualification, the compiler expands functions in line, propagates constant arguments, passes arguments in registers, and monitors module-level static variables. Use the following `-Qoption` keywords to refine these interprocedural optimizations:

`ip_args_in_regs=0`

Disables the passing of arguments in registers. By default, external functions can pass arguments in registers when called locally. Normally, only static functions can pass arguments in registers, provided the address of the function is not taken and the function does not use a variable number of arguments.

`ip_inline_max_calls=N`

Used with the inline heuristics (`-Qinl_heur`) option. This option changes the default number of call-sites to inline. Note that `N` call-sites are inlined only if that many call-sites meet the minimum inline criteria. For more information, see the [“Using In-line Heuristics \(-Qinl_heur n\)”](#) and [“Criteria for In-Line Function Expansion”](#) sections.

`ip_inline_max_stats=n`

Sets the allowable number of intermediate language statements and expressions for a function that is expanded in line. The number `n` is a positive integer. The number of intermediate language statements usually exceeds the actual number of source language statements. The default is set to the maximum number of 230.

`ip_inline_max_total_stats=n` Sets the maximum increase in the number of intermediate language statements and expressions for each function that is expanded in line. The number `n` is a positive integer. By default, each function can increase to a maximum of 2000 statements.

Analyzing the Effects of Multifile IPO (-Qipo_c, -Qipo_S)

The `-Qipo_c` and `-Qipo_S` options are useful for analyzing the effects of multifile IPO, or when experimenting with multifile IPO between modules that do not make up a complete program.

Use the `-Qipo_c` option to optimize across files and produce an object file. This option performs optimizations as described for `-Qipo`, but stops prior to the final link stage, leaving an optimized object file. The default name for this file is `ipo_out.obj`. You can use the `-Fe` option to specify a different name as shown in the following example:

```
prompt> icl -G6 -Qipo_c -Fe filename a.cpp b.cpp c.cpp
```

Use the `-Qipo_S` option to optimize across files and produce an assembly file. This option performs optimizations as described for `-Qipo`, but stops prior to the final link stage, leaving an optimized assembly file. The default name for this file is `ipo_out.asm`. You can use the `-Fe` option to specify a different name as shown in the following example:

```
prompt> icl -G6 -Qipo_S -Fe filename a.cpp b.cpp c.cpp
```

Using In-line Heuristics (-Qinl_heur n)

Use the `-Qinl_heur n` option to inline the most frequently used call-sites. When you specify this option, call-sites are treated as leaf routines. For more information on in-lining and the minimum in-lining criteria, see [“Controlling Inline Expansion of User Functions \(-Obn, -Qip no inlining\)” in Chapter 5](#).

The `-Qinl_heur n` option only takes effect when used together with the `-Qip` or `-Qipo` option. The value of `n` can be 0, 1, 2, or 3.



NOTE. *Inline heuristics 1 through 3 are all profile-guided. If no profile information is available the compiler uses the default inline heuristics (that is, `n = 0`) and issue a warning. See [“Profile-Guided Optimization \(PGO\): Three Phases” in Chapter 5](#) for profile information procedures.*

There are four variations on the `-Qinl_heur n` option:

- `-Qinl_heur 0` Default in-line heuristics. These heuristics work both with or without available profile information.
- `-Qinl_heur 1` In-line the `N` most frequently executed call-sites.
- `-Qinl_heur 2` In-line the `N` most frequently executed call-sites. After each in-lining, the execution counts are scaled and the list of frequently executed call-sites is updated.
- `-Qinl_heur 3` In-line the `N` most frequently executed call-sites scaled by size of the callee. After each in-lining, the execution counts are scaled and the list of call-sites is updated.

By default, `N = 100`. You can change this value by using the option `-Qoption,c,-ip_inline_max_calls=N`. Note that `N` call-sites will be inlined only if that many call-sites meet the minimum inline criteria.

Criteria for In-Line Function Expansion

For a routine to be considered for in-lining, it has to meet certain minimum criteria. There are criteria to be met by the call-site, the caller, and the callee. The call-site is the site of the call to the function that might be in-lined. The caller is the function that contains the call-site. The callee is the function being called that might be in-lined.

Minimum call-site criteria:

- The number of actual arguments must match the number of formal arguments of the callee.

- The number of return values must be the same as the callees' number.
- The data types of the actual and formal arguments must be compatible.
- No multi-lingual in-lining is allowed. Caller and callee must be written in the same source language.

Minimum criteria for the caller:

- At most 2000 intermediate statements will be inlined into the caller from all the call-sites being inlined into the caller. You can change this value by specifying the option
`-Qoption,c,-ip_inline_max_total_stats= new value`
- The function must be called or have its address used if it is declared as `static`. Otherwise, it will be deleted.

Minimum criteria for the callee:

- Is not considered infrequent due to the name. Routines that contain the following substrings in their names are not inlined: `abort`, `alloca`, `denied`, `err`, `exit`, `fail`, `fatal`, `fault`, `halt`, `init`, `interrupt`, `invalid`, `quit`, `rare`, `stop`, `timeout`, `trace`, `trap`, and `warn`.

Once these criteria are met, the compiler picks the routines whose inline expansions provide the greatest benefit to program performance. This is done using the following default heuristics. When you use profile-guided optimizations, a number of other heuristics are used (see [“Profile-Guided Optimization \(PGO\): Three Phases” in Chapter 5](#) and [“Using In-line Heuristics \(-Qinl heur n\)”](#) earlier in this chapter for more information).

- The default heuristic focuses on call-sites in loops or calls to functions containing loops.
- When profile information is available, the focus changes to the most frequently executed call-sites. Also, the default inline heuristic does not allow the in-lining of functions with more than 230 intermediate statements, or the number specified by the option
`-Qoption,c,-ip_inline_max_stats`.
- The default inline heuristic stops when it detects direct recursion.
- The default heuristic will always inline very small functions that meet the minimum inline criteria. By default, functions with 15 or fewer intermediate statements are inlined. This limit can be modified with the option `-Qoption,c,-ip_inline_min_stats`.

Index

Symbols

!= operator 8-4
#assert preprocessor directive 7-3, 8-3
#define preprocessor directive 7-3
#include preprocessor directive 3-3
#unassert preprocessor directive 8-3
#undef preprocessor directive 7-3
.asm extension 2-3, 6-2
.i extension 2-3, 7-2
.lib extension 2-3
.obj extension 2-3, 6-4
; Semicolon character 8-4
< operator 8-4
<= operator 8-4
-? option 10-1
_ _cplusplus predefined macro 7-4
_ _DATE_ _ macro 8-5
_ _FILE_ _ macro 8-5
_ _ICL predefined macro 7-4
_ _LINE_ _ macro 8-5
_ _STDC_ _ macro 8-5
_ _TIME_ _ macro 8-5
_ asm keyword 8-2
_CHAR_UNSIGNED predefined macro 7-5
_CPPUNWIND predefined macro 7-5
_DLL predefined macro 7-5

_M_IX86=n predefined macro 7-5
_MSC_VER=900 predefined macro 7-4
_MT predefined macro 7-5
_WIN32 predefined macro 7-4
} Closing brace 8-4
} Right brace character 8-4

A

a.exe file 6-5
Addition of 0 elimination 4-8
Alignment constraints 2-15, 12-1
ANSI standard 8-1, 8-2
 Compiler limits 8-1
 Conformance to 8-1
 Extensions 8-1, 8-3, 8-4
Application development 1-2
ARGSUSED symbol 10-4
Arguments
 In registers B-2
asm keyword 8-2, 8-4
Assembly files 2-3

B

Bit fields 8-3
 Types 8-3

C

- C language dialects 8-2
 - Extended ANSI 8-1, 8-2
 - Strict ANSI 8-2
 - Strict ANSI conformance 8-1
- C option 7-2
- c option 6-4
- Casting
 - Of integer constants 8-3
 - Of pointers 8-3
- Catastrophic messages 10-4
- Comma character
 - In enum list 8-4
- Command-line
 - Multiple filenames 2-2
 - Options (syntax) 2-2
 - See also Compilation 2-1
 - Syntax 2-1
- Comparisons
 - Of pointers to types 8-4
- Compilation
 - Command-line syntax 2-1, 2-2
 - From MSVC++ 2-2
 - Passes 2-15
 - Source processing
 - See Source processing 7-1
- Compilation and execution differences
 - Intel vs Microsoft 9-4
- Compilation phases
 - c (suppress linking) 6-4
 - E (preprocess only) 7-2
 - EP (preprocess and omit #line directives) 7-2
 - P (preprocess only) 7-2
 - S (assembly code listing) 6-2
 - Zs (syntax check only) 6-2
- Compiler
 - Default behavior 2-15
 - Front-end program 3-4, 3-5
 - Invocation from MSVC++ 2-2
 - Invocation from the command-line 2-1

- Compiler options
 - Quick guide 2-4
- Compiler pragmas
 - Limitations on 9-1
- Configuration file 3-2
- Conformance to ANSI 8-1
- Conformance to standards 8-1
- Constant folding optimization 4-8

D

- D option 7-1, 7-3
- Debugging
 - Preparing for 6-6
- Debugging, symbolic 6-6
- Default libraries 11-2
- Diagnostic messages 10-2
 - Catastrophic 10-4
 - Command-line 10-2
 - Errors 10-3
 - Remarks 10-3
 - Suppressing of 10-4
- Division by 1 elimination 4-8
- double data type 8-3
- Double floating-point precision 4-7

E

- E option 7-1, 7-2
- Empty declarations 8-4
- Empty file 8-2
- Entry point 3-5
- enum data type
 - Base types 8-3
 - Tag name 8-2
- Environment variables 3-1
- EP option 7-2
- Error messages 10-3
 - Displaying 10-4
- Errors, internal 10-3

Executable output 6-4
Extended ANSI 8-1
Extended floating-point precision 4-7

F

-Fa option 6-2, 6-5
-Fe option 6-2, 6-5
Filename extensions 2-3
 .asm files 6-2
 .i files 2-3
 .lib files 2-3
 .obj files 2-3, 6-4
 for C++ 2-3
Files
 Assembly code listing 6-2
 Input 2-2
 Object 6-4
 Output 6-5
 Preprocessing 7-2
Floating-point
 Conformance 4-8
 Data type 8-3
 Double precision 4-7
 Order performed 4-8
 Precision 4-7
Floating-point arithmetic precision 4-7
Floating-point arithmetic precision per
 architecture 4-8, 4-9
-Fo option 6-2, 6-5
Function calls 5-2
Function Order List
 Example 5-10
 for Profile Guided Optimizations 5-9
 Guidelines 5-10

H

-help option 10-1

I

IEEE 754 4-7
INCLUDE variable 3-3
Initialization
 Pointer 8-3
 Static array 8-2
 struct 8-2
 union 8-2
In-line assembly language
 Insertion 8-4
In-line function expansion
 Criteria for B-4
In-line heuristics B-3
Input files 2-2
Integral types 8-3
Intermediate language statements B-2, B-3
Internal errors 10-3
Interprocedural Optimization
 multifile example 5-3
 Multiple Files 5-2
Interprocedural optimization B-2
Invocation
 From MSVC++ 2-2
 From the command-line 2-1
 See also Compilation 2-1
Invocation syntax 2-1
-ip_inline_max_calls option B-4
-ip_inline_max_stats option B-5
IPO Multifile Executable
 Using a makefile 5-4

K

Keywords
 __asm 8-2
 asm 8-2, 8-4

L

Labels 8-4

- Language conformance 8-1
- LIB variable 3-1
- libm_chk.lib file 11-3
- Libraries
 - Default 11-2
 - Managing 11-1
 - Math libraries 11-3
 - Routines, in-lining of 4-6
- Library files 11-2
 - Use of 11-1, 11-2
- Library functions
 - Expansion of 4-6
- Library routines
 - In-lining of 4-6, B-5
- link (linker) program 3-4
- Linker
 - Direct support of 3-5
- lint program 10-4
- lint-specific comments 10-4
 - ARGSUSED 10-4
 - NOTREACHED 10-4
 - VARARGS 10-4
- long float data type
 - See double data type 8-3
- Loop interchange optimization 4-8, B-2
- lvalue type 8-3
- _MSC_VER=900 7-4
- _MT 7-5
- _WIN32 7-4
- Makefile
 - To create a Multifile IPO 5-4
- Makefile dependencies
 - Printing 7-6
- Managing libraries 11-1
- masm (assembler) program 3-4
- Math libraries 11-3
- Math library file 11-3
- Memory
 - Optimizations B-2
- Messages
 - Diagnostic 10-2
 - Format 10-2
 - Remarks 10-3
 - Suppressing of 10-4
- Microsoft Visual C++ 32-bit edition for Windows iv
- MMX intrinsics 13-1
 - Compare intrinsics 13-9
 - General support intrinsics 13-4
 - Logical intrinsics 13-9
 - Packed arithmetic intrinsics 13-5
 - Shift intrinsics 13-7
- Multiplication by 1 elimination 4-8

M

- Macros, predefined
 - __cplusplus 7-4
 - __DATE__ 8-5
 - __FILE__ 8-5
 - __LINE__ 8-5
 - __STDC__ 8-5
 - __TIME__ 8-5
 - __ICL 7-4
 - _CHAR_UNSIGNED 7-5
 - _CPPUNWIND 7-5
 - _DLL 7-5
 - _M_IX86=n 7-5

N

- nologo option 10-1
- NOTREACHED symbol 10-4

O

- O1 option 4-1
- O2 option 4-1
- Object files 6-4
 - Generation of 6-4
- Od option 4-1
- Oi option 4-6

- Oi- option 4-6, 4-8
- Op option 4-7, 4-8
- Op- option 4-7
- Optimizations B-2
 - Cloning B-1
 - Constant folding 4-8
 - Floating-point precision 4-7
 - In-line expansion B-1
 - Interprocedural B-2
 - Loop interchange 4-8, B-1, B-2
 - Memory B-2
 - Perform no optimizations 4-1
 - Procedural 5-2, 5-3, B-3
 - Qualifying (-QW0) B-1
 - Types of 4-2
- Options 2-15
 - ? (print a summary list of icl options) 10-1
 - C (preserve comments in preprocessed source output) 7-2
 - c (suppress linking) 6-4
 - D (defining macro) 7-3
 - E (preprocessing output to stdout) 7-2
 - EP (preprocess and omit #line directives) 7-2
 - Fa (assembler output filename) 6-5
 - Fe (executable output filename) 6-5
 - Fo (object output filename) 6-5
 - help (print a summary list of icl options) 10-1
 - nologo (disable sign-on message) 10-1
 - O1 (intraprocedural optimization) 4-1
 - O2 (intraprocedural optimization) 4-1
 - Od (perform no optimizations) 4-1
 - Oi- (disables in-lining of libraries) 4-6, 4-8
 - Oi (enables in-lining of libraries) 4-6
 - Op (favors conformance to floating-point standards over optimization) 4-7
 - Op- (favors optimization over conformance floating point standards) 4-7
 - Oy (enable use of ebp register) 6-6
 - P (preprocessing output to file) 7-2
- Passing to other tools 3-4
- Project-specific 3-2
- QA (asserting names) 7-3
- QA- (suppress predefined macros and assertions) 7-3
- QH (include file pathnames) 7-5
- QIfdiv- (disable floating-point division check) 11-4
- QIfdiv (enable floating-point division check) 11-3
- Qinl_heur <n> (use in-line heuristics) B-3
- Qip (interprocedural optimization) 5-2, B-2
- Qipo (enable interprocedural optimization between files) 5-3
- Qipo_c (optimize across files and produce an object file) B-3
- Qipo_S (optimize across files and produce an assembly file) B-3
- QM (print makefile dependencies) 7-6
- Qnobss_init (allocate zero-initialized variables) 12-2
- Qpc (floating-point precision per architecture) 4-8, 4-9
- Qprec (improve floating-point precision) 4-8
- Qprof_dir 5-10
- Qprof_gen 5-7
- Qprof_genx 5-9
- Qprof_use 5-7
- Quse_asm (use assembly file) 6-4
- QW (passing options) 3-4
- QW0 (qualifying optimizations) B-1
- Qwd (disable soft diagnostic) 10-5
- Qwe (change severity of soft diagnostic to error) 10-5
- Qwn (limit number of errors reported) 10-6
- Qwr (change severity of soft diagnostic to remark) 10-5
- Qww (change severity of soft diagnostic to warning) 10-5
- S (assembly code listing) 6-2
- To change defaults 3-2
- u (suppress predefined macros and assertions) 7-3
- U (undefining names) 7-3
- Unsupported (Microsoft) 9-3

- W (suppress warnings) 10-4
- w (suppress warnings) 10-4
- Za (strict ANSI) 4-8, 8-2
- Zi (symbol table for debugging) 6-6
- Zp (specify alignment) 12-1
- Zs (syntax check only) 6-2
- Output files
 - Naming of 6-5
- Overflow 8-4
- Oy option 6-6

P

- P option 7-1, 7-2
- Passing options
 - To other tools 3-4
- PATH variable 3-1
- Pathnames 7-5
- PCH Files
 - Differences in Support 9-4
- Pointers 8-3
 - Assignment to interchangeable types 8-4
 - Comparing to difference types 8-4
 - To integers 8-3
- pp-number syntax 8-3
- Pragmas
 - Limitations on 9-1
- Precision
 - Floating-point arithmetic 4-7
 - Floating-point arithmetic per architecture 4-8, 4-9
- Precompiled Header Files
 - Differences in Support 9-4
- Predefined macros 8-4, 8-5
- Predicate names
 - Defining 8-3
- Preprocessor directives 7-3, 8-4
- Procedural optimization 4-1, 5-2, 5-3, B-3
 - Multifile 5-3
- Profile data
 - explicit dump 5-12

- Profile Guided Optimization
 - Utilites 5-11
- Profile Guided Optimizations
 - Function Order List 5-9
- profmerge utility 5-11
- proforder utility 5-11
- Publications
 - See related publications iv

Q

- QA option 7-1, 7-3
- QA- option 7-1, 7-3
- QH option 7-5
- QIfdiv option 11-3
- QIfdiv- option 11-4
- Qinl_heur <n> option B-3
- Qip option 5-2, B-2
- Qipo option 5-3
- Qipo_c option B-3
- Qipo_S option B-3
- QM option 7-6
- Qpc and -Qrct
 - Alternatives to 4-9
- Qpc option 4-8, 4-9
- Qprec option 4-8
- Qprof_dir 5-10
- Qprof_gen 5-7
- Qprof_genx 5-9
- Qprof_use 5-7
- Qrcd and -Qrct
 - Rounding Options 4-10
- Quse_asm option 6-4
- QW option 3-4, 3-5
- QW0 option B-1, B-2
- Qwd option 10-5
- Qwe option 10-5
- Qwn option 10-6

-Qwr option 10-5
-Qww option 10-5

R

Related publications iv
Remark messages 10-3
Remarks
 Enabling 10-4
Response file 3-2

S

-S option 6-2
Source processing 7-1
 Defining macros (-D) 7-3
 Preparing for debugging (-Zi) 6-6
 Undefining names (-U) 7-3
Static functions B-2
Static variables B-2
stdout 7-2
Strict ANSI conformance 8-1, 8-2
Strictest alignment constraint 2-15
struct data type 8-2
 Members 8-2, 8-3
Structures 8-2
 Alignment of 12-1
Subtraction of 0 elimination 4-8
Symbolic debugging 6-6
 Optimizations for 6-7
 Support for 6-7
Symbols
 Asserting 7-3
 Defining 7-3
 Undefining 7-3

T

Tag declarations 8-4
Target architecture iv

TMP variable 3-1
Tools you need 1-1

U

-U option 7-1, 7-3
-u option 7-3
union data type
 Members 8-3
Unions
 Alignment of 12-1

V

VARARGS symbol 10-4

W

-W option 10-4
-w option 10-4
Warning messages
 Suppression of 10-4
Warning messages Suppressing 10-4

Z

-Za option 4-8, 8-2
-Zi option 6-6
-Zp option 12-1
-Zs option 6-2